

Towards an improved measurement of the electron electric dipole moment

Henry Ashworth

**Submitted for the degree of Doctor of Philosophy of the
University of London.**

Imperial College

July, 2008

Abstract

This thesis builds upon previous generations of an experiment to measure the electron electric dipole moment (EDM) undertaken at the Centre of Cold Matter, Imperial College. The electron EDM is of great interest as an improved measurement will constrain the parameterisation of various extensions to the standard model, and a non-zero measurement would be direct evidence of T-symmetry violation.

The method used to probe the EDM was an implementation of a separated oscillatory field molecular interferometer based upon a pulsed supersonic beam of cold YbF. The electron EDM (d_e) was exposed through analysis of the splitting between magnetic sublevels due to both Zeeman and EDM interactions. Carefully controlled switching of applied magnetic and electric fields allowed the independent measurement of the EDM interaction contribution, and in turn the calculation of d_e .

Principally this thesis is concerned with the understanding of systematic effects that have recently become apparent. Two main contributions are discussed here: effects arising from the curvature of the electric field and effects associated with the spatial extent of the molecular pulse. Resolution of these systematic effects primarily involved improved control and understanding of the electromagnetic environment. Firstly the experiment was adapted to utilise a single electric field using one set high precision electric field plates. These were then used to construct a radio frequency (rf) transmission line, replacing the more traditional rf loop scheme.

To better understand electromagnetic fields which may affect the EDM measurement, techniques were subsequently developed to measure components of the electric, magnetic and rf fields throughout the experiment.

A small data set was acquired to ensure that the modifications functioned as expected. This data was taken over a period of several months, during which time the experiment was never fully optimised. This data set achieved a sensitivity of $2.15 \times 10^{-27} \text{ e.cm}$ (68 % bootstrapped confidence interval).

Contents

1	Introduction	11
1.1	Electric dipole moments	11
1.1.1	The EDM in particle physics	13
1.1.2	Origins of electric dipole moments	15
1.2	Measuring an EDM	15
1.2.1	EDM experiments	17
1.2.2	The separated oscillatory field technique	24
1.3	Experiment theory	26
2	Implementation	30
2.0.1	Infrastructure	30
2.0.2	The YbF beam	31
2.1	Lasers	33
2.1.1	Ablation laser	33
2.1.2	Pump - probe laser	34
2.2	Electric fields	34
2.2.1	Power supplies and control	35
2.2.2	High precision field plates	36
2.3	Radiofrequency fields	37
2.3.1	Radiofrequency structures	38
2.3.2	Radiofrequency mixers, synthesiser and amplifier	39
2.4	Static magnetic field	41
2.4.1	Magnetic field coils	41
2.4.2	Magnetic shielding	42
2.5	Instrumentation	43
2.5.1	Photomultiplier tubes and laser induced fluorescence	43
2.5.2	Magnetometry	44
2.5.3	Data acquisition	44

3	Experimental Method	48
3.1	The molecular beam	48
3.1.1	The time of flight (TOF) data	49
3.2	Scanning parameter spaces	50
3.3	State preparation	51
3.4	A single resonance rf experiment	51
3.5	A double resonance rf experiment	55
3.6	EDM Experiment aspects	57
3.6.1	Measuring the EDM	57
3.6.2	Analysis of the EDM measurement	59
3.6.3	Development of the EDM measurement scheme	61
3.7	Noise	63
3.7.1	Detector noise	63
3.7.2	Source noise	65
3.7.3	Magnetic field noise	65
3.7.4	Switching pattern noise rejection	68
3.7.5	Phase locking and mains noise	69
4	Systematic effects	71
4.1	Review of the EDM measurement concept	71
4.1.1	Probing for possible systematic effects	72
4.2	Specific systematic effects	73
4.2.1	Systematic effects arising from the use of a multiple electric field regions . .	73
4.2.2	Electric field asymmetry systematic	78
4.2.3	Investigation of analysis channels across the time of flight	79
4.3	Systematic effects conclusion	86
5	Measurement of electric, magnetic and rf fields.	88
5.1	Field mapping using a single rf resonance	88
5.1.1	Scanning a single rf transition	88
5.1.2	Mapping the electric field	91
5.1.3	Mapping the magnetic field	94
5.1.4	Mapping the rf field	96
5.2	Probing the electric field using a Ramsey interferometer	99
5.2.1	Dynamics of the separated oscillatory field approach	99
5.2.2	Ramsey interferometer results	101

6 EDM Results	103
6.1 EDM data analysis	103
6.1.1 Bootstrap of EDM data before veto	104
6.1.2 Magnetic field veto	105
6.2 Discussion of results	107
7 Conclusion	112
7.1 Main thesis results	112
7.2 Future progress	113
7.2.1 Short term development	113
7.2.2 Long term development	117
7.3 Outlook	119
A Experiment parameters	120
A.1 Molecular beam optimisation	120
B Theoretical description of rf pulses	122
C YbF Stark shift	126
D Experiment operation guide EDM Data	127
D.1 Preparation of the experiment	127
D.2 Taking EDM Data	129
E Photographs of the EDM experiment	131

List of Figures

1.1	EDM/spin combinations	13
1.2	Comparison of predicted electron EDM values	14
1.3	Feynman diagram	15
1.4	Ramsey's neutron interferometer	18
1.5	Polarisation of YbF	22
1.6	Motional magnetic field systematic configuration	23
1.7	Interferometer scheme illustration	25
1.8	Interferometer sequence illustration	26
1.9	Simplified YbF structure	27
1.10	Interferometer output illustration	28
2.1	Illustration of experimental apparatus	31
2.2	Electric field control schematic	35
2.3	Electric field plate assembly	37
2.4	Transmission line rf configuration	39
2.5	Radiofrequency pulse synthesis hardware	40
2.6	Magnetic field coils	42
2.7	EDM experiment data structures	45
2.8	Computer hardware interface	47
3.1	Example time of flight (TOF)	50
3.2	Example parameter scan	50
3.3	Radiofrequency frequency scan (rf plates)	53
3.4	Radiofrequency power scan (rf plates)	54
3.5	Radiofrequency power scan (rf loops)	54
3.6	Interference scan data	56
3.7	Eight point EDM measurement scheme	58
3.8	Simplified interferometer lineshape	59
3.9	Example switching waveforms	62
3.10	Simple EDM noise analysis	64

3.11	Power density spectrum of source noise	65
3.12	Power density spectrum of magnetic field noise	66
3.13	Bootstrap example	67
3.14	EDM data veto analysis	68
3.15	Power density spectra of switching patterns	69
3.16	Moving-window power density spectrum of magnetic mains noise	69
4.1	Electric field between guard and centre regions under asymmetric applied potentials	74
4.2	Experimental evidence supporting the electric field gradient systematic	77
4.3	Systematic effect arising from B - and E -shifts	78
4.4	Comparative time binned analysis of SIG and CAL channels using rf loops	80
4.5	Comparative time binned analysis of EDM data using loops and plates	80
4.6	Comparative time binned analysis of EDM data B channel using loops and plates	81
4.7	Single loop magnetic field	82
4.8	Block level B channel analysis unsigned by B state	83
4.9	B -channel analysis unsigned by B -state	84
4.10	Fit of interferometer phase shift to varying beam alignment	85
4.11	Inferring molecule position by velocity class	85
5.1	Simplified single resonance rf experiment	89
5.2	Example of raw scan over rf transitions	90
5.3	Fit of rf scan to rf transition lineshape	90
5.4	Scan analysis for mapping	91
5.5	Electric field map, binned, with error bars	93
5.6	Scan over Zeeman split rf transition	94
5.7	Binned magnetic field map	95
5.8	Plot of raw magnetic field	95
5.9	Raw magnetic field maps, applied field and ambient field	96
5.10	Scan of rf amplitude showing Rabi flopping	97
5.11	rf amplitude map	97
5.12	rf mapping systematic error	98
5.13	Ramsey interferometer fringes	101
5.14	Ramsey fringes for each state of the electric field	101
6.1	Raw block data	103
6.2	Raw EDM data analysis	104
6.3	Bootstrapped confidence intervals, before the magnetic field veto	105
6.4	EDM data veto investigation	106

6.5	EDM data veto by time	107
6.6	EDM data analysis - with magnetic field veto	107
6.7	Bootstrapped confidence intervals, after the magnetic field veto	108
6.8	Comparison of vetoed and non-vetoed bootstrap distributions	109
6.9	Comparison of bootstrapped vetoed and rejected data distributions	110
7.1	Future EDM experiment development	113
C.1	YbF Stark shift	126
E.1	Photographs of EDM experiment	131

List of Tables

1.1	Theoretical calculations of atomic electron EDM enhancements	19
1.2	Effective electric fields for polar molecules and molecular ions.	22
1.3	Overview of key EDM results	24
3.1	Analysis of stepped interferometer scheme data.	60
4.1	Field plate voltages for asymmetric electric field systematic	75
4.2	Field plate voltages for asymmetric electric field test experiment	76
6.1	Symmetric and absolute Bootstrapped EDM confidence intervals, before magnetic field noise veto.	105
6.2	Bootstrapped EDM confidence intervals, after veto.	107
6.3	Sensitivities of the EDM experiment, after magnetic field veto.	108
A.1	Main EDM experiment parameters	121

Acknowledgements

There are many people to thank. Ed, Ben, Jony and Mike must come at the top of the list. For the past few years they have ensured that the EDM experiment an incredibly stimulating place to be. Beyond everything they have taught me in the lab, this thesis has benefitted immeasurably from their suggestions and review. Ed and Ben have always found time to discuss physics, no matter how busy. Jony has taught me not just physics, but instilled some degree of style into my mathematics and programming.

I'd like to thank my EDM peers, Paul and Dhiren, with whom I shared the coveted task of changing the laser dye, as well as the Centre for Cold Matter group as a whole for numerous coffee break debates. Jon Dyne deserves special mention for his excellent engineering assistance, as does Sanja for helping me navigate the challenges of student registration.

My family have been a great source of support, offering particularly favourable rents. I recall a previous tutor once explaining that a good physicist should really be happy to fix their own bike: I thank my parents for buying those bikes in the first place. And allowing me to disassembly - often terminally - everything from cars to the office computer.

Finally I should like to thank my Tom and Gemma, my long suffering housemates, who quietly tolerated substantial electricity bills as programs ran through the night. In addition, I thank Phil, Kate and Claire for their continual encouragement. I'm sorry to disappoint them all by not *actually* 'reversing time'.

Though I may be the first in my family privileged to study for a Doctorate in Physics, I'm certain that I'm not the first who would of liked to. With that thought I dedicate this thesis to the memory of my Grandfather, Ted, whom I am certain would have made a fantastic scientist given my opportunities.

Chapter 1

Introduction

This thesis describes recent progress made at the Centre for Cold Matter, Imperial College in measuring the permanent electric dipole moment (EDM) of the electron. This first chapter begins by examining the importance of EDMs, then goes on to look at several EDM measurements. Finally I introduce the basic principles of our experiment.

1.1 Electric dipole moments

Permanent electric dipole moments of elementary particles offer a very important insight into particle physics. Various experiments searching for EDMs of the proton [1], neutron [2], electron [3] and muon [4]¹ all find results consistent with zero. However, popular extensions to the Standard Model predict non-zero EDMs to lie close to current experimental limits. At present these limits place severe constraints upon particle physics models. A further increase in sensitivity would further restrict, or even rejection the validity these models. Secondly, the discovery of a non-zero EDM would directly imply violation of T-symmetry.

Symmetry violation is an important concept in modern physics, and can be traced right back to the Dirac equation. This describes relativistic quantum mechanical behaviour. However, the equation suggests a series of solutions of negative energy. In the context of electrons, these were originally interpreted as ‘holes’ in a negative sea of electrons. Several years later in 1932 the positron was discovered - the first example of antimatter. Immediately this raised the question of why the universe seemed dominated by matter. Unhappy with the argument that somehow the initial conditions of the universe were asymmetric; Pauli argued that Dirac’s hole-theory must be incorrect, and insisted that the Dirac equation treats matter and antimatter symmetrically[6].

In the 1950’s theorists began to probe the nature of symmetry more directly. Parity (P) symmetry² was first investigated when theoretical investigations by Lee and Yang concluded strong

¹The muon EDM was inferred by McNabb from analysis of the year 2000 dataset of the muon g-2 experiment [5].

²A parity transformation is the reversal of one spatial coordinate.

and electromagnetic forces seem invariant under parity symmetry, though they believed certain weak interactions violate the symmetry. Parity violation was shortly after observed in the β -decay of Cobalt 60 by Wu et al. [7] and independently by Garwin et al. [8]. Soon it was realised that the product symmetry CP was conserved in these experiments³. However, in 1964 CP-violation in kaon decay was demonstrated by Christensen et al. [9].

Today, the product of CP symmetry with T-symmetry, CPT-symmetry, is however thought to be a ‘good’ symmetry, with the corollary of this is that CP violation implies T violation. However, given that violations of C, P, CP symmetries have been observed, it may seem shortsighted to rely too heavily on CPT conservation. In fact CPT symmetry has a solid theoretical basis, relating it to Lorentz symmetry⁴, stating that local field theories which have Lorentz symmetry must also be CPT symmetric. Such theories encompass the Standard Model and SUSY, and hence this physics is believed to be CPT symmetric.

T-symmetry implies that the laws governing a process are symmetric under the transformation $t \longrightarrow -t$. Commonly this is described as T-reversal, but really there is no suggestion that time is reversed. Certain examples in classical physics are T-symmetric. In the classical world, we ought expect it equally common to find spinning tops spinning clockwise, as it is to find them spinning anticlockwise. However, in general complex *macroscopic* systems are not T-symmetry invariant due to thermodynamic effects. Baking a cake, cracking an egg, or lighting a match can only happen in one direction as, under the second law of thermodynamics, the entropy of systems tend to increase. To convince yourself of this you need only look at a video recording of such an event; it should be immediately obvious whether the recording is playing backwards or forwards. To avoid these effects, we need to consider a simpler, *microscopic* system, ideally a single fundamental particle. The problem then, is finding a microscopic system with a symmetry violating macroscopic property.

A permanent electric dipole moment is a measure of the displacement of the centre of charge from the centre of mass. Naïvly, imagine we alter an electron, such that it’s charge distribution is egg shaped, the charge displacement represented by the bulge of the egg. We know electrons have spin, so let us imagine the ‘egg’ is rotating about the major axis. Then we expect to find four types of eggs: either spinning clockwise, or anticlockwise, bulge up, bulge down. We can take the two eggs spinning clockwise and argue these are simply the two spinning anti-clockwise, under the transformation $t \longrightarrow -t$. If EDM is zero, (i.e. no bulge - a spherical egg), then we expect to see one type of egg.

In practice, we only ever see one type of electron, where spin up and spin down are related by $t \longrightarrow -t$. Our experiences of the universe suggest strongly that the electron has a single degree of

³C-symmetry means that the laws of physics are symmetric under the charge conjugation transformation. Though this is described by specific transformations in quantum field theory, it can be more intuitively understood by considering invariance under the substitution of all charges $q \longrightarrow -q$ in classical electromagnetism.

⁴Broadly, Lorentz symmetry states that under a Lorentz transformation (a rotation or boost), that the laws of physics remain unaffected

freedom - there is just one type, spin up or spin down. The best example perhaps is the periodic table, the mechanics of which necessitate the electron to have one degree of freedom - spin.

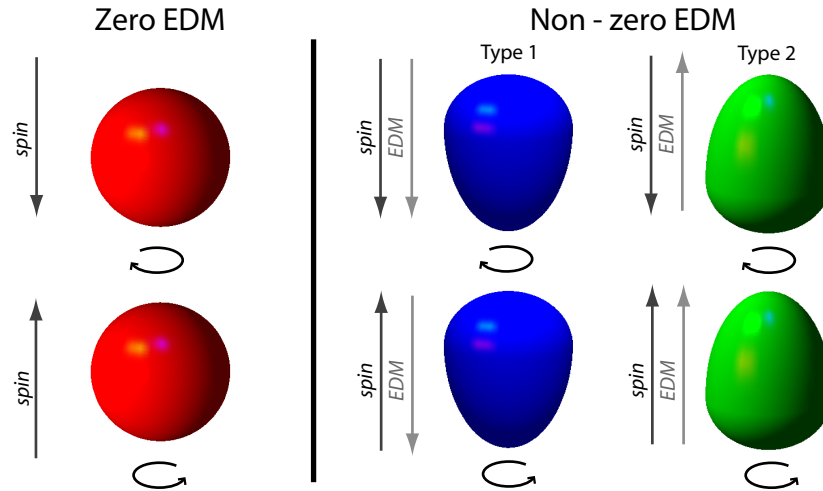


Figure 1.1: In the case that the electron EDM is zero (red), we see one type of electron, in spin up and spin down configurations. The introduction of an EDM leads to 2 types of electron, each spin up or down, with ‘egg up’ or ‘egg down’. This is at odds with our experience of the world, hence we conclude a non-zero EDM implies that T-symmetry must be violated.

Given this, if we measure a non-zero EDM, it must be the case that the dipole moment and spin vector prefer to be either parallel or anti-parallel. This preference necessitates is T-symmetry violation. If the EDM is zero, then T-symmetry is not violated. In section 1.1.1 I will introduce the concepts that can lead to a permanent EDM.

This can succinctly be expressed considering that the EDM of a particle introduces the energy shift $\Delta E = \vec{d}_e \cdot \vec{E}$. The vector \vec{d}_e , the EDM, has an expectation proportional to $\langle \phi | \vec{J} | \phi \rangle$, where J is the angular momentum of the electron. Under T reversal, $\vec{J} \longrightarrow -\vec{J}$. For a state that is invariant under T-reversal, the EDM must be zero. Accordingly, a non-zero EDM implies T-symmetry violation.

1.1.1 The EDM in particle physics

The Standard Model (SM) was devised in the early 1970’s as a concise description of particles consistent with both quantum mechanics and special relativity. However, it has several limitations: it can describe the electromagnetic, strong and weak forces, but not gravity; it does not explain baryogenesis, the dominance of matter over antimatter [6]; and it cannot solve the Hierarchy problem, why the mass of the Higgs Boson is so much smaller than the Planck mass.

Recently several ideas have been proposed to address these issues, the most popular of which is supersymmetry (SUSY). This partially solves the hierarchy problem, allowing the Higgs Boson to have a small mass⁵. It also allows for more CP violation to explain baryogenesis. As such, SUSY is a very important step towards greater understanding of fundamental physics. The problem

⁵Why the Higgs boson is so light is unaddressed by SUSY, and known as the μ problem.

comes in testing the theory. Direct search measurements, which aim to produce all of the physics described by SUSY are extremely complex and expensive, and as such take many years to develop.

Experiments to measure an electron EDM play an important rôle. They provide a narrow window deep into high energy physics. The standard model predicts very small EDMs of order 10^{-40} e.cm[10]. This is in fact far too small to for either the current generation, or any proposed experiments to measure. However, extensions to the standard model predict much larger EDMs, of order $10^{-24} - 10^{-30}$ e.cm. The range each model spans is due to the particular parameterisation chosen. With the most sensitive experiment to date probing to a sensitivity of 8×10^{-28} e.cm, it is certainly possible to make measurement of this order.

As the uncertainty in these EDM measurements is gradually reduced there are two possible outcomes. If with improving accuracy the EDM is consistent with zero (a null result), then the models becoming increasingly constrained, with the implication that at some point certain models may be rejected. If however a non-zero EDM is measured, then it provides a data point around which the models may be calibrated. Additional, as previously discussed, a non-zero EDM is an important result in its own right as an example of T-symmetry violation.

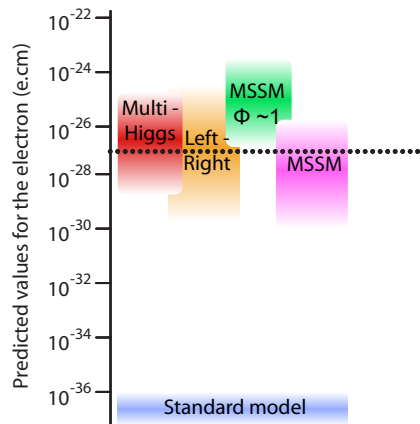


Figure 1.2: Predicted values for the electron EDM, along with the current experimental limit (dotted line).

In general, CP- violation is expressed theoretically as a complex phase relation between two fundamental fields. The standard model accommodates the CP- violation in the K_0 system by allowing a single complex phase, δ , in the Cabibbo-Kobayashi-Maskawa (CKM) Matrix. Fortson[11] notes, first order W-quark interactions come in pairs for which this phase cancels, so the contribution to any EDM is small (see figure 1.3.b).

Supersymmetry doubles the number of particles by allowing each particle a *superpartner*. This increase in the number of particles allows for many more interactions, and with them more allowable CP- violation. Fortson[11] gives the example of spin zero bosons, such as the electron's partner the selectron, which can have CP- violating interactions with electrons and quarks. An example of such an interaction is compared in figure 1.3 to a CP- violation the standard model.




Figure 1.3: Feynman diagrams of boson - quark interactions in both SUSY and the standard model. *a*: In SUSY, the two phases which control emission and reabsorption may be different. *b*: In the standard model the limited number of interactions, constrains the phase to be the same for emission and reabsorption.

1.1.2 Origins of electric dipole moments

First of all let's recall what an electric dipole moment is. Imagine a simple system of two point charges $+q$ and $-q$, separated by a distance \mathbf{r} , the electric dipole moment is defined as $\mathbf{p} = q\mathbf{r}$. This vector \mathbf{p} points from the negative to the positive charge, and gives us an idea of the 'polarity' of the system. In measuring the permanent electric dipole moment of the electron we are really trying to infer something about the structure of the charge distribution. A single spherically symmetric charge implies zero EDM, so an obvious question is perhaps why the electron, which is well known to be a point charge, should have anything other than a zero EDM.

This leads to the origins of the EDM. Fermions, electrons in the case of this thesis, are not static, rather they continuously emit and absorb bosons. These particles may become polarised by their interaction with the bare fermion. The electron may therefore be viewed not as just a single point particle, but as a cloud of charge. If the vacuum polarisation then gives a displacement to the positive and negative charges the electron will have a permanent distortion to the charge distribution - an EDM. In the standard model, there is only one phase which allows for CP-violation, so the emission and subsequent reabsorption of a boson must cancel. This process renders the charge distribution of the fermion symmetric. However, in SUSY, the plethora of new particles allows for processes with different phases. These do not cancel fully, and can create an asymmetric charge distributions about the electron, which we may observe as an EDM.

1.2 Measuring an EDM

There are various systems where we can look for EDMs. The most extensive EDM experiments have been carried out on bare particles, atoms and molecules, searching for EDMs of electrons, neutrons, protons, and muons.

The most general principle in measuring an EDM is to investigate how the EDM interaction $d_e \cdot \vec{E}$ modifies the energy of a system. This is somewhat complicated by the fact that as the EDM

and Zeeman effect both couple to the electron spin, they are of a similar mathematical form; we are looking for an energy shift under applied \vec{E} and \vec{B} , $h\nu = 2|\mu_B B + d_e E|$. Additionally, the Zeeman effect is obviously known to be many orders of magnitude stronger, so limited control of the magnetic field will likely obscure any evidence of an EDM. Clearly, to resolve the two effects a method for controlling magnetic fields is needed.

So, how small is the EDM? The units of e.cm are however a little opaque. Let us anticipate the size of spectral feature we are searching for. Firstly, assume that the electron EDM lies just below the current limit; $d_e \leq 6 \times 10^{-28}$ e.cm say. For YbF at a reasonably attainable field of 20 kVcm^{-1} , this leads to $\Delta U [\text{eV}] = 1.7 \times 10^{10} d_e [\text{e.cm}]$, so we expect a detuning of order $\Delta U = 1.0 \times 10^{-17} \text{ eV}$, a frequency of 2.5 mHz. Resolving such a small energy shift presents significant technical challenges that form the basis for this thesis⁶.

For electrons and protons, it is not practical to use the bare particle. The sensitivity of such experiments is limited, as the application of an electric field would accelerate the particle into the electrodes⁷, so, in the case of the electron for a more precise measurement we look to larger *neutral* systems: atoms and molecules. The mechanism by which the electron imparts its EDM to the atom or molecule is not immediately obvious. A first guess might be that as atoms contain electrons, then the effect of the EDM might be detectable through small perturbations to the gross atomic or molecular spectra.

In the search for the electron EDM, the first systems to be examined were atoms. A series of theoretical results were derived investigating the possibility of measuring an electron EDM in atomic systems. Initially, the electrons and nucleus the atoms were considered to be point particles. Schiff showed that in an external electric field the nucleus and electrons reconfigure, such that the net average electric field on any of the particles averages to zero, masking any electron EDM to the observer. However, he went on to show that if the constituents are allowed a finite size, this screening is incomplete. The EDM interaction is attenuated by a factor of order 10^4 , but is not zero. In diamagnetic atoms, the finite size of the nucleus leads to a Schiff moment. In this case EDM is a result of the misalignments of the nuclear charge and dipole moment distribution. In paramagnetic atoms the measured EDM is due largely to an unpaired spin.

A second way to escape Schiff screening is through the relativistic dynamics of the electrons. Sandars went on to show that for high Z (proton number) paramagnetic atoms the EDM shielding scales with the nuclear charge of the atom $Z^3 \alpha^2$ [13]. This implies, despite the Schiff screening, there can be considerable enhancements, where Z is large enough.

Finally, we should consider the neutron EDM. Neutrons, having no charge, may be probed

⁶It is impossible to express the EDM in such units without making some assumptions about the experiment used. Without wishing to introduce new concepts, I note the numbers presented here are based on our EDM experiment. This uses an interferometer with a coherence time of approximately 1000 μs , and assumes a molecular polarisation of 0.68. These values are introduced later in this chapter.

⁷In fact, an early, simple electron EDM experiment was carried out in this manner by Nelson ([12]), yielding the bound $|d_e| \leq 10^{-13} \text{ e.cm}$. In addition the muon EDM was measured in a storage ring [4].

directly in spin precession experiments. A discussion of neutron EDM experiments is presented in section 1.2.1.

Atomic EDM experiments suffer from two key limitations: the magnitude electric field that can be applied and systematic effects. Though it is possible to create large electric fields in the lab (1000 kVcm^{-1}), maintaining a uniform field over a large area, without breakdown or leakage currents is extremely challenging⁸. However, it was noted by Sandars et al. ([15]), that heavy polar molecules have strong internal ‘effective’ fields. Using polar molecules leads to 4 or more orders of magnitude increase the size of the EDM interaction, for typically attainable fields. This can be attributed to fact the polar molecules can be polarised to a higher degree than atoms. Furthermore, using polar molecules helps to reject certain systematic effects caused by the motion of the molecule in the electric field⁹.

1.2.1 EDM experiments

Before discussing our scheme to measure the electron EDM, I will briefly discuss other important EDM experiments.

Neutron EDM experiments

Neutron EDM experiments were first carried out in the 1950’s by Ramsey [17], using his ‘separated oscillatory field’ technique¹⁰. This approach (detailed in section 1.2.2) offers an extremely sensitive technique for measuring relative energies using the precession frequencies of spin different states as a probe. The experiment is based upon a beam of neutrons. Initially unpolarised neutrons are emitted from a pile. The first step then is to polarise them using a polished magnetised mirror. This prepares all¹¹ neutrons in the same state. The neutrons now enter into a region of homogeneous magnetic and electric field - the interaction region. The next step is to prepare a beam of neutrons in a coherent superposition of spin states using an rf pulse. The theory supporting this operation is developed in appendix B. As the neutrons, continue to propagate through this region, the opposite spin states precess with opposite angular frequencies, at a rate governed by the state energies. As the neutrons come to the end of the interaction region the spin states have accrued a relative phase, dependent on the magnetic and electric dipole moment interactions. This phase is then measured through the application of a second rf pulse. The amplitude of a given spin state after this pulse is dependant on the phase accrued in the interferometer. In practice to measure the state amplitudes a second magnetic analyser is used to select a certain polarisation, before the neutrons pass into a counter. The energy shift which causes this relative phase accrual depends on both

⁸Currents originating from the electric field plates are an awkward source of magnetic fields, which must be dealt with carefully (See Condylis [14] 6.2.4). In practice leakage and discharge currents limit attainable fields and increase field switching times.

⁹See Hudson’s thesis ([16].3) for a discussions of motional magnetic field effects, which are suppressed using polar molecules.

¹⁰This approach is often called the ‘Ramsey interferometer’.

¹¹In fact, Ramsey’s polariser was approximately 85% efficient.

EDM and Zeeman interactions, so through the application of different combinations of electric and magnetic fields, the EDM can be inferred. Ramsey's experiment constrained the neutron EDM to $d_n < 5 \times 10^{-20} \text{ e.cm}$ [17].




Figure 1.4: Ramsey's 1957 neutron beam EDM experiment can be considered a prototype for the YbF electron EDM experiment detailed in this thesis. Unpolarised neutrons enter from the left. After being polarised (prepared) using a magnetic mirror, they pass into the interaction region. The first rf pulse creates a superposition of spin states. Under applied electric and magnetic fields, these accrue a relative phase, ϕ_I . A second rf pulse rotates the superposition, such that the amplitude of either state can be used to infer ϕ_I . With suitable combinations of \vec{E} and \vec{B} , the magnitude of the EDM interaction can be inferred from ϕ_I .

The latest experiment carried out at the Institut Laue-Langevin (ILL) have further constrained the limit of the neutron EDM to $d_n < 3.0 \times 10^{-26} \text{ e.cm}$ [2]. In this experiment ultra cold neutrons pass through thin magnetised foil, which creates a sample of neutrons with aligned spins. The neutrons are stored in large cell, under applied \vec{E} and \vec{B} fields. Whilst in storage the neutrons undergo a separated oscillatory field experiment, similar in nature to Ramsey's neutron beam experiment [17]. Two second pulses of oscillatory magnetic field were applied at a frequency close to the Larmor frequency, separated by 130 seconds of free evolution. The final step of the experiment was to count the population of N_\uparrow and N_\downarrow neutrons. Key improvements over Ramsey's experiment include increased flux ($\sim 100\times$), and trapping the neutrons which increases the integration time from milliseconds to minutes. Additionally, the neutron storage vessel contains ^{199}Hg , which acts as a magnetometer¹².

Though the standard model admits a small neutron EDM through the violation of CP -symmetry of order $d_n = 10^{-32} - 10^{-33} \text{ e.cm}$, many extensions to the standard model predict much larger neutron EDMs of order $10^{-25} - 10^{-26} \text{ e.cm}$. Clearly the ILL experiment is already heavily constraining this region [18].

The next generation of ultra cold neutron experiment, proposed by the scientists at University of Sussex, promises a higher electric field, longer storage times and more neutrons. The neutrons are to be made by down scattering a neutron beam in 0.5 K liquid helium. This will lead to a factor of 50 improvement in neutron density. Liquid helium is also a better electrical insulator than the attainable lab vacuum, so a gain of 4-6 in electric field will be possible. The experiment is

¹²The presence of a second species, acting as a magnetometer within the sample is termed a comagnetometer. The advantage to this approach is that, rather than perhaps in cells surrounding the experiment, the magnetic field over *exactly* the same region as the neutron sample is monitored.

Atom		α
Rb	[21]	24
Cs	[21]	125
Au	[22]	260
Tl	[19]	-585
Fr	[22]	910

Table 1.1: Theoretical calculations of atomic electron EDM enhancements

designed such that the neutrons do not leave the He, which should lead to much longer coherence times. Overall, a factor of 100 improvement in sensitivity is expected.

Atomic EDM experiments

Modern electron EDM measurements utilise atoms or molecules. The calculation of electron EDM enhancements are far from trivial, and well beyond the scope of this thesis([18],[19],[20]). Such analysis encapsulates the complexities of calculating relativistic wavefunctions, Schiff screening and the appropriate enhancement. For an atomic experiment we introduce an enhancement α to relate observed net EDM to that of one electron.

$$\hat{H}_{EDM} = -\alpha d_e \hat{F} \cdot \vec{E}_{\text{ext}} \quad (1.1)$$

This Hamiltonian describes the interaction of the EDM d_e , along the angular momentum axis \hat{F} , in the applied field \vec{E}_{ext} . Typical enhancements (α) are given in table 1.1. These vary in scale considerably. This is due in part to the fact the enhancement scales cubically with proton number, so the heaviest nuclei give rise to the largest enhancements.

The current leading measurement of the electron EDM was made by Commins et al. using a thallium atomic beam [3]. The choice of thallium¹³ is a compromise between maximising the enhancement, and practicality (Francium is highly radioactive). The result of Commins' experiment was $d_e = (6.9 \pm 7.4) \times 10^{-28} \text{ e.cm}$, which implies $|d_e| \leq 1.6 \times 10^{-27} \text{ e.cm}$.

The first generation of this experiment originally implemented two beams of thallium. The experiment was qualitatively similar to Ramsey's 1950 neutron EDM experiment in its use of separated oscillatory field spectroscopy. Considering a single beam, the experimental procedure was to prepare the thallium in the $6P_{1/2}, F = 1, m_F = 0$ state, then use an rf transition ($\pi/2$ -pulse) then created a superposition of $F = 1, m_F = \pm 1$ states. As with Ramsey's experiment magnetic and electric fields were applied, during which time the superposition evolved, dependent on the Zeeman and EDM interactions. The final phase of the superposition is measured by driving a second identical rf transition, and probing the residual $F = 1, m_F = \pm 1$ population. The second thallium beam was used with the opposite velocity, and served to reject a large systematic effect

¹³Thallium is also different to the other elements quoted, as it has a negative enhancement. This is due to the fact the EDM is manifest due to mixing of the 6s and 6p wavefunctions close to the nucleus. The structure of thallium is $(6s^2)(6p)$, which results in the removal of an electron from a filled shell, leaving a hole. This in turn leads to a change in sign of the enhancement[23].

due to the motion of the molecules through the magnetic field. For this generation of experiment the electron EDM was constrained to $|d_e| \leq 4 \times 10^{-27} \text{ e.cm}$ [24].

The second generation of the experiment implemented a comagnetometer, as well as common mode magnetic field rejection. This was achieved using two pairs of dual species beams, bringing the electron EDM to the limit quoted above. Despite extensive systematic rejection the experiment was eventually limited by systematic effects and the size of energy shift that could be resolved.

Francium experiment

From table 1.1, it is clear that in terms of the enhancement factor, francium, the heaviest alkali metal, is an excellent candidate for an electron EDM experiment. However, this the experiment will be somewhat complicated by the fact that francium has a half life of between 4 and 20 minutes, so, the atoms must be produced on demand. At present a collaboration of groups are working towards trapping and producing Francium [25]. Furthermore, to reduce the number of atoms required, the EDM experiment will take place in a magneto-optical trap (MOT) built on a chip. This experiment is still in development.

^{199}Hg experiment

Diamagnetic EDM experiments probe for an EDM of proportional to the Schiff moment of the nucleus. To first order, the electrons screen any external field. However, taking into account the finite size of the nucleus, there is a small residual P- T- odd component of the electric field - the Schiff moment.

The most sensitive experiment to measure such EDM was performed by Fortson's Group [26]. He constrained the EDM of ^{199}Hg to be $d_{^{199}\text{Hg}} < 2.10 \times 10^{-28} \text{ e.cm}$. As in electron EDM experiments, no T-violation was observed. In order to relate this result to the underlying particle theory, the Schiff moment of the ^{199}Hg nucleus must be calculated [27]. This can then be more directly compared to CP-violating nucleon-nucleon interactions, which in turn can constrain QCD phase angles.

The experiment used two cells filled with mercury vapour. Combinations of magnetic and electric fields were used to measure any EDM interaction $\mathbf{d}_{\text{Hg}} \cdot \mathbf{E}$ by observing changes in the nuclear Zeeman spin precession frequency. Each cell experienced the same magnetic field, but with a reversed electric field vector. Using pump-probe spectroscopy the difference in precession rate of each cell was measured, which would be due to the Hg EDM. Over a period of 30 sec, the mercury cells were pumped with circularly polarised light strobed at the Larmor frequency. This built up a spin polarisation aligned to the pump light in the cell. During the probe phase, the light was attenuated and plane polarised. The spins now aligned to the applied magnetic field, and modulated the transmitted probe light at the Larmor frequency.

A modification of this experiment uses 4 Hg filled cells [28]. In this case, two of the cells have no applied electric field, and serve as magnetometers, which should reject magnetic field noise.

The other important nuclear EDM experiment was performed using thallium fluoride by Sanders and Hinds [29] at Oxford, then continuing at Yale. This experiment primarily probed the proton EDM, as the thallium nucleus has an unpaired proton. Again Ramsey’s separated oscillatory field technique was used to measure the small energy shift arising from the proton EDM. The result of the experiment was that $d_p = -3.7 \pm 6.3 \times 10^{-23} \text{ e.cm}$ [29].

In summary, over the past 50 years, experiments have measured EDMs consistent with zero in protons, electrons, or bare neutrons. The current limit of the electron EDM has already constrained some phases within extensions to the standard model, but various model parameterisations still leave several orders of magnitude where we may find an EDM. Therefore, in support of these theories, continued efforts to measure any EDM with increased sensitivity would provide valuable information to particle physics.

Polar molecule electron EDM experiments

In this section I will discuss the motivation behind our EDM experiment - an ytterbium fluoride separated oscillatory field interferometer.

Commins’ Thallium EDM experiment stands as the most sensitive measurement of the electron EDM. However, despite a long period of development, systematics eventually slowed progress after a factor of 4 improvement in sensitivity. Our approach, using YbF, has the advantage that YbF has a greater overall enhancement factor than Tl.

In an atomic system the EDM enhancement is greatest in heavy paramagnetic atoms. Close to the nucleus there are strong relativistic effects[30], leading to an enhancement proportional to Z^3 . In order to measure the EDM however, the atom must be polarised. Sushkov et al [31] noted that polar molecules could be used as sensitive probes of the electron EDM, as in a molecular system, the atoms are strongly polarised along the internuclear axis. This averages to zero unless the molecule is polarised with a modest electric field. The overall effect is that using polar molecules, very large ‘effective fields’ (E_{eff}) are attainable. Kozlov has calculated the effective field for YbF to be $2.5 \times 10^{10} \text{ Vm}^{-1}$ [32], whilst the largest homogeneous¹⁴ field we can make in the lab is of order 10^4 Vcm^{-1} . So, now we must relate the polarisation of the molecule to the energy shift due to the EDM,

$$\Delta U = -\eta d_e E_{\text{eff}} \tag{1.2}$$

The term η is the polarisation of the molecules which varies with the applied electric field, as shown in figure 1.5.

¹⁴Though it is possible to create larger inhomogeneous fields, they are not generally suitable for EDM experiments.




Figure 1.5: Polarisation of the YbF molecule within an electric field. The extent to which the molecule aligns to the field limits the sensitivity to the EDM.

Molecule		$E_{\text{eff.}}$ (GV/cm)
YbF	[32]	26
BaF	[20]	7.4
HgF	[33]	99
PbF	[33]	-29
HfH ⁺	[34]	-17
PtH ⁺	[34]	73

Table 1.2: Effective electric fields for polar molecules and molecular ions.

From figure 1.5, using typical attainable lab fields of 20 kV cm^{-1} we can achieve molecular polarisations of approximately 70%. Combined with the effective field above (equation 1.2), the basic interaction is 600 times stronger than the TI experiment.

There are a range of other molecules we could consider. For measuring electron EDM interactions we need a heavy nucleus, and an unpaired electron. Some of the molecules for which effective fields have been calculated are shown in table 1.2.

The second reason for using polar molecules is to avoid a systematic effect arising from motion induced magnetic fields. This was first discussed in Sandars 1964 measurement of the EDM of the Cesium atom [35], and is discussed in greater depth by Lamoreaux [36]. An atom moving through an electric field experiences a magnetic field, $\vec{B}^{\text{mot}} = \vec{E} \times \frac{\vec{v}}{c^2}$. With typical beam velocities of 450 ms^{-1} this appears as a field aligned to the x -axis, $B_x^{\text{mot}} = 30 \text{ nT}$, which reverses with E . In atomic systems, the axis of quantisation is not necessarily orthogonal to x . Therefore, an atom will generally sample a fraction of \vec{B}_x^{mot} , such that the splitting is dependent upon the field $B' = |\vec{B}^{\text{mot}}| \sin(\theta) + |\vec{B}_z| \cos(\theta)$. The angle θ is the rotation of the atoms axis of quantisation about the y -axis. In this case the Zeeman interaction becomes dependant upon the electric field, making it indistinguishable from the EDM interaction. Using a polar molecules in an idealised system the molecules are strongly aligned to the z axis, such that the effect of \vec{B}_x^{mot} is heavily suppressed.




Figure 1.6: Vector diagram supporting the explanation of the motional magnetic field. The atoms travelling with velocity \vec{v} in the $+y$ direction experience a magnetic field aligned to B_x caused by the propagation through the electric field. As the atoms are not strongly aligned the z -axis the sample some fraction of this field.

Using PbO to measure the electron EDM

DeMille is in the process of using PbO to make an improved electron EDM measurement [37], predicting a sensitivity of $|d_e| < 10^{-31}$ e.cm. PbO is a paramagnetic molecule, with a similar effective electric field to YbF. However, DeMille is able to use a vapour cell, rather than molecular beam, resulting in a much denser molecular sample may be established. This is practical due to the physical nature of PbO; it may readily be vaporised inside the cell and is thermodynamically stable. Additionally the $a(1)$ state, where the experiment is to take place is easily populated using laser excitation and has an extremely high polarisability. Modest electric fields of 10 Vcm^{-1} generate polarisations $P \simeq 1$. In comparison, YbF would need around 100000 Vcm^{-1} to achieve similar polarisations. At this voltage a vapour cell implementation would be impossible. Though a proof of concept experiment has been undertaken [38], an EDM measurement has not yet been published.

Using HfH⁺ ions to measure the electron EDM

From table 1.2, it is clear that the molecular ions HfH⁺ and PtH⁺ have large effective fields. Cornell is working on a molecular ion EDM measurement, which may measure the electron EDM to $|d_e| < 6 \times 10^{-29}$ e.cm [39]. Using a molecular ion brings several advantages. Firstly, like PbO, they are easy to polarise using small electric fields. Secondly, the ions may be trapped, leading to long coherence time experiments. There are no suggestions of working with PtH⁺.

Using PbF to measure the electron EDM

Schafer-Ray [40] recently calculated that using PbF an EDM measurement could be performed with little sensitivity to the magnetic field. He found that when performed on certain states, if the magnetic fields were shielded to $1.0 \mu\text{G}$, then the magnetic dipole moment contribution to the electron EDM is below 2.0×10^{-32} e.cm. Though this makes PbF an interesting candidate, a viable experiment would require much spectroscopy, detector and source development.

A summary of all leading EDM measurements is given in table 1.3.

EDM Probed	system	result (e.cm)	date	author	[ref.]
neutron	n	$d_n \leq 5 \times 10^{-20}$	1957	Ramsey	[17]
		$d_n \leq 3 \times 10^{-26}$	2006	Baker (ILL)	[2]
proton	TlF	$d_p = -3.7 \pm 6.3 \cdot 10^{-23}$	1989	Hinds	[1]
muon	μ	$d_\mu < 2.8 \times 10^{-19}$	2004	McNabb	[4]
^{199}Hg	^{199}Hg	$d_{^{199}\text{Hg}} \leq 2.10 \times 10^{-28}$	2001	Fortson	[26]
electron	e	$d_e \leq 10^{-13}$	1959	Nelson	[12]
	Tl	$d_e \leq 1.6 \times 10^{-27}$	2002	Commins	[3]
	YbF	$d_p = 0.2 \pm 3.2 \cdot 10^{-26}$	2001	Hudson	[41]
	PbO	-	-	DeMille	[42]
	HfH ⁺	-	-	Cornell	[39]
	PbF	-	-	Schafer Ray	[40]

Table 1.3: Overview of EDM experiment results discussed in this review.

1.2.2 The separated oscillatory field technique

Throughout the progression of EDM experiments the problem of resolving such a small energy shift was solved using a separated oscillatory field technique. The first example of this can be traced back to Ramsey’s neutron EDM experiment [17].

In any of the systems discussed, a (non-zero) electron EDM will lead to an energy shift of certain states under applied electric field. Analogous to the Zeeman effect, the interaction splits degenerate states depending upon the angular momentum projection. Taking the pair of hyperfine sub levels $F = 1, m_F = \pm 1$, and using equation 1.2 (for an atomic system), the EDM leads to the splitting

$$\Delta U_{\text{EDM}} = \alpha d_e \vec{E}_{\text{ext}} \cdot [\vec{F} - (-\vec{F})] = -2\alpha d_e \vec{E}_{\text{ext}} \cdot |\vec{F}| \quad (1.3)$$

To measure this energy difference, a superposition of these hyperfine sublevels is created. Ignoring the Zeeman effect, under an electric field, the oppositely signed sublevels begin to accrue relative phases due solely to the EDM interaction. This free evolution is due to the dynamics of the time independent Schrödinger equation.

$$i\hbar \frac{d}{dt} |\psi_n(t)\rangle = E_n |\psi_n(t)\rangle \quad (1.4)$$

$$|\psi_n(t)\rangle = e^{\frac{iE_n t}{\hbar}} |\psi_n(0)\rangle \quad (1.5)$$

By comparing these phases we can perform a fantastically sensitive measurement of the electron EDM. In order to measure this phase we have to create a quantum interferometer to coherently compare these two states.

The process is analogous to an optical interferometer, where a beam splitter is used to creates an optical superposition which can be used to measure small distances (changes in path length). An outline of the interferometer process is shown in figure 1.7. Specific details on the experimental implementation, as well as the spectroscopy of YbF follow in the next section.

The YbF EDM experiment is based upon the two hyperfine levels, $F = 0$ and $F = 1$. To start with the system is prepared in the state $|F = 0\rangle$.

The next step is to use a radiofrequency magnetic field to drive σ_{\pm} transitions, creating a coherent superposition of the hyperfine magnetic sublevels. Labelling the hyperfine states $|F, m_F\rangle$, $|0, 0\rangle \rightarrow 1/\sqrt{2}(|1, 1\rangle + |1, -1\rangle)$ (figure 1.7.b). This is similar to the first beam splitter in an optical interferometer.

Magnetic and electric fields are now applied for approximately 1000 μs . This splits the degeneracy of the $|1, \pm 1\rangle$ states. The complex amplitudes of each state now evolve freely. The Stark shift has no overall effect: the $m_F = \pm 1$ sublevels are shifted identically, hence evolve no relative phase. However the EDM and Zeeman interactions both lead to a relative phase difference between the states. The analogous situation in the case of an optical interferometer, is to introduce a phase shift by changing a path length in one arm.

The next step is to compare the phase difference of the states. A second rf pulse is applied which transfers some amplitude back to the $|F = 0\rangle$ state, dependant upon the relative phase. The final amplitudes of both the $|F = 0\rangle$ and $|F = 1\rangle$ states are dependant upon the phase accrued whilst the molecules were in a coherent superposition. This amplitude varies sinusoidally with increasing phase difference (a full derivation is given in section [?]). The final step is to measure the population of either $|F = 0\rangle$ or $|F = 1\rangle$.

A summary of the experiment is shown pictorially in figure 1.8. Only the experimental elements key to the separated oscillatory field interferometer are shown. The molecules travel left to right, passing through the sequence of events described in 1.7 above.

The approach taken here implements a superposition of magnetic sublevels which are immune to Stark shift inducing phases, but sensitive to both Zeeman and EDM effects. Through application of suitable electric and magnetic field combinations we can separate the two, resolving the EDM phase, and hence the EDM, d_e .




Figure 1.7: Interferometer scheme. The molecular initially consists of with molecules thermally populated into in both $F=0$ and $F=1$ states. The first action is to remove molecules from $F=1$, which we do using a 553 nm transition to the next electronic state, from where they largely decay to dark states. The first ‘beam splitter’ creates a superposition of magnetic sublevels using a π -pulse. After a period of free evolution a second π -pulse is applied before the population of the $F=1$ state is probed.




Figure 1.8: Illustrated summary of interferometer experiment. This representation of the apparatus used the electron EDM experiment is similar to many previous beam experiments. The key features are a region of applied magnetic/electric field placed between two beam splitter rotations, which in this case are driven using rf transitions.

1.3 Experiment theory

Having set out a basic outline, I will describe in detail the interferometer, deriving the lineshape as well as considering a detection scheme.

A thorough description of the structure of YbF is beyond the scope of this thesis. For a discussion of the spectroscopy of YbF, refer to earlier work on YbF [43]. What follows is merely sufficient to explain our technique.

The electronic shell configuration of Yb is $\text{Yb}([\text{Xe}](4f)^{14}(6s)^2)$. To build YbF, a single 6s electron is transferred to the $\text{F}([\text{He}](2s)^2(2p)^5)$ atom. The resulting YbF radical is ionically bound with this single valence electron. Of all the states of the molecule, perhaps the most important for us is the electronic, vibrational and rotational ground state $X^2\Sigma^+(\nu = 0, N = 0)$, as this is where we perform the experiment. The fluorine (nuclear spin 1/2) couples with the valence electron resulting in singlet $F=0$ and triplet $F=1$ hyperfine levels, split by approximately 170 MHz (see appendix C). These hyperfine levels form the interferometer, and are used throughout figure 1.7. The overall splitting of the $F=0$ and $F=1$ states is sensitive to the applied electric field through the Stark effect. However the relative splitting of the $|1, 1\rangle$ and $|1, -1\rangle$ states is sensitive only to Zeeman and EDM interactions.

The scheme described previously needs a little development. Starting with the preparation phase we need to ensure that all molecules are in the $X^2\Sigma^+(\nu = 0, N = 0), F = 0$ state. Spectroscopically, higher rotational and vibrational states play no part in the experiment. However, our YbF source has finite intensity, so as molecules in excited states are essentially lost, we do not wish to populate these states from the outset. By cooling the YbF via supersonic expansion to around 5 K we are able to obtain a large population of the ground state. Condylis calculates that at 10 K the probability of a molecules being in the ground vibrational state is 100%, the ground rotational state 35%¹⁵. Cold molecules will populate both $F=0$ and $F=1$, so we require a preparation step to remove any population from the $F=1$ levels. To do this we use the 553 nm

¹⁵By way of comparison, Condylis notes using an effusive oven source at 1500 K the probability of being in the ground vibrational state is 35%; the probability of being in the ground rotational state only with only 0.023%

$Q(0)F=1$ resonance $X^2\Sigma^+(\nu=0, N=0) F=0 \longrightarrow A^2\Pi_{1/2}(\nu=0, N=0) F=1$ to drive molecules to the next electronic state (see figure 1.9). From there they largely decay to ‘dark states’, which are which are no longer experimentally significant¹⁶. For the most part in this thesis I will describe the YbF system in terms of the $|F=0\rangle$ and $|F=1, m_F=\pm 1\rangle$, assuming that the electronic, rotational and vibrational state to be the $X^2\Sigma^+(\nu=0, N=0)$ ground state.

The next step is preparation of the coherent superposition - the first beam splitter in analogy to the optical interferometer.




Figure 1.9: Simplified YbF structure. Shown are the levels relevant to the EDM experiment. Particularly important is the 553 nm $Q(0)F=1$ transition, used for pumping and probing.

To drive the transition $|F=0\rangle \longrightarrow \frac{1}{\sqrt{2}}(|1, -1\rangle + |1, 1\rangle)$ we use an rf pulse. The $F=0$ and $F=1$ levels are separated by ~ 170 MHz. The rf is linearly polarised along y -axis. This is achieved (in the first instance) using a loop in the $x-z$ plane. The rf is applied on resonance, for a duration t_{rf} , at a power of P_{rf} . These parameters must be carefully set such that one-half period Rabi oscillation is performed - a π -pulse¹⁷. A full derivation of the rf transition is provided in appendix B.

The interferometer is sensitive to both Zeeman and EDM interactions. The total measured phase, (ϕ_I) , is the sum of $\phi_{de} = 2d_e\eta E_{\text{eff}} \frac{t}{\hbar}$ and $\phi_B = 2\mu_B B \frac{t}{\hbar}$.

The complete interferometer process can be described mathematically in five steps. The derivation of the π -pulse is described in appendix B.

¹⁶We can in fact re-pump from the $X^2\Sigma^+(\nu=0, N=2)$ state, to the $X^2\Sigma^+(\nu=0, N=0) F=0$ state. This has the effect of recovering some of the lost population from pumping the ground state, $F=1$ level. This has not been attempted within this thesis, though it was attempted by Hudson, and will likely be used in our final dataset [16].

¹⁷Ramsey’s original experiment upon a two level system $|1\rangle, |0\rangle$ was driven to the superposition $(|1\rangle + |0\rangle)/\sqrt{2}$. This is a rotation of $\pi/2$ upon the Bloch sphere, so was termed a $\pi/2$ -pulse. In our system we use twice the power, driving our molecules into the (degenerate) $|1, \pm 1\rangle$ states, so our pulses are termed π -pulses.

- (a) Prepared State: $|0, 0\rangle$
- (b) First π -pulse: $\frac{1}{\sqrt{2}} (|1, -1\rangle + |1, 1\rangle)$
- (c) Free evolution: $\frac{1}{\sqrt{2}} (|1, -1\rangle e^{-i(\phi_{de} + \phi_B)} + |1, 1\rangle e^{i(\phi_{de} + \phi_B)})$
- (d) Second π -pulse: $|0, 0\rangle \cos(\phi_{de} + \phi_B) + \frac{e^{-i}}{\sqrt{2}} \sin(\phi_{de} + \phi_B) [-|1, -1\rangle + |1, 1\rangle]$
- (e) Probe $|1, \pm 1\rangle$ states $|\langle 1, -1|1, -1\rangle|^2 + |\langle 1, 1|1, 1\rangle|^2 = \sin^2(\phi_{de} + \phi_B)$

To measure the interferometer output phase a second ‘recombining’ π -pulse is used. As shown in equation *d* above, this leads to amplitude being split between the $|F = 0\rangle$ and $|F = 1\rangle$ states, as a function of the total interferometer phase. The resulting amplitude of the $F=0$ and $F=1$ hyperfine states oscillate sinusoidally. These are interference fringes, conceptually identical to fringes seen in a classical optical interferometer.

The final task is to measure the population of either the $F=0$ or $F=1$ state. We use a process called ‘Laser induced fluorescence’ (LIF). this involves probing the $|F = 1\rangle$ state with the same laser source used for state preparation. The probe laser drives a $Q(0)F=1$ transition to the $A^2\Pi_{1/2}$ electronic state. We use a photomultiplier tube (PMT) to detect the photons emitted from the spontaneous decay of this state to infer the interferometer output.

As seen in equation 1.3.e, the final measured probability amplitude of the $F=1$ state is of the form $\sin^2(\phi_I)$. If over many individual interferometer experiments we continuously increment the magnetic field amplitude, we trace out an interference curve (figure 1.10). The EDM interaction would appear identical, although given the scale of the electron EDM, the pitch is very large - if not infinite!




Figure 1.10: Interferometer output illustration. The two curves represent the interferometer output, under the reversal of the applied electric field (with greatly a exaggerated EDM interaction strength). The EDM is then inferred from the relative phase shift of the curves.

Using this technique we can only measure the total phase due to the Zeeman and EDM interactions. As both effects are functionally similar we must rely on manipulating the electric and magnetic fields to resolve terms. The most obvious solution is to shield external fields to a level where the Zeeman shift phase is much less than that due to the EDM. Then, by recording a set of experiments at different electric fields we might see a change in count rate due to the EDM. Comparing phase terms, for and electron EDM of order 10^{-28} e.cm (just below the current ex-

perimental limit), and typical lab fields of kVcm^{-1} , we would need to reduce internal magnetic fields to below 15 fT. Given a typical lab magnetic field of $50 \mu\text{T}$, even a shielding factor of 1000 would be woefully insufficient. Trimming the magnetic field is a second option, but again would be experimentally challenging at this level.

Instead we opt for a scheme whereby we step the electric and magnetic fields. Knowing the EDM to be small, we step the magnetic field such that we step between $\pm\phi_I \simeq \phi_B = \pi/4$. We also step the applied electric field, $E = \pm 12 \text{ kVcm}^{-1}$. In the simplest case this leads to an interferometer phase associated with 4 states, $\{B_+E_+, B_+E_-, B_-E_+, B_-E_-\}$. The EDM phase is then isolated by examining the difference in interferometer phase in states, $1/4(B_+E_+ + B_-E_+ - B_+E_- - B_-E_-)$. This scheme rejects DC magnetic fields and with sufficient samples will reject AC magnetic field noise.

The sensitivity of the experiment is discussed in chapter 3, where various sources of noise are considered. An important motivation in the experiment however is that the precision of the final measurement scales as $1/\sqrt{N}$, where N is the number of individual measurements we make. As $N \propto \text{time}$, it is clear that no significant gains in sensitivity may be made through a longer data run: if it takes one month to measure to a level of 10^{-27} e.cm , it will take more than 8 years integrate down to 10^{-28} e.cm .

In chapter 2 the apparatus used is discussed, whilst a more detailed description of our implementation is presented in 3.

Chapter 2

Implementation

In the previous chapter, the basic interferometry scheme used to measure the electron EDM was explained. In this chapter I discuss the experimental apparatus, though very detailed implementation detail is omitted, unless it is required to support later chapters. This thesis contains no detailed CAD or schematics, as it would add little to the reader’s understanding. Much of the apparatus is in fact unchanged since the work of Condylis [14], and in some cases Hudson [16] and Redgrave [44]. For brevity, where appropriate I refer to their theses.

From the brief description of the experiment given in the previous chapter, some of the key elements have already been introduced. The experiment implements a form of molecular interferometer based upon a pulsed beam of YbF. Beam splitter transitions are performed using short rf pulses, whilst a 553 nm laser is used for pumping (state preparation) and detection. Electric and magnetic fields are applied whilst the molecules are in a quantum superposition, isolating the EDM interaction. All of these systems are explained in detail in this chapter. Additionally, I introduce data acquisition, the vacuum assembly, magnetic shielding and instrumentation.

2.0.1 Infrastructure

The heart of the experiment comprises of a large cylindrical vacuum chamber 20 cm x 2 m. The chamber is orientated vertically, as shown in figure 2.1. It is divided into two regions by a gate valve and skimmer, with the gate valve open under normal conditions. The lower chamber, used for creating the molecular beam, contains the source assembly. This section is opened frequently, around once a month, primarily to clean the target. The upper chamber houses the rest of the experiment (detectors, field plates, B-coils, etc). This chamber remains sealed permanently in order to protect the electric field plates¹. The vacuum is maintained by four turbo pumps, two large pumps on the lower chamber², two smaller pumps on the upper chamber. To allow for independent

¹After opening the upper chamber it is often the case that the field plates need to be re-conditioned (see §2.2.2) to run at high voltage.

²The upgrade from one to two large turbo pumps on the lower chamber was made recently, with a view to running at 50 Hz. In practice, no net gain in EDM sensitivity (per root unit time) was observed, so the choice was taken to continue to operate at 25 Hz.




Figure 2.1: Illustrative rendering of EDM experiment. Shown here are the orientations of the key experimental components. The definition of axes given here will be used throughout this thesis.

servicing of both chamber regions the upper and lower pumps are backed independently. During typical operation the chamber pressures are 10^{-4} mBar and 5×10^{-7} mBar for the lower and upper chambers respectively.

The convention in this thesis is that the y -axis is defined by the direction of propagation of the beam, and the z -axis is defined by the electric field vector, assumed ordinarily to be homogeneous throughout the interaction region.

2.0.2 The YbF beam

The experiment is based upon a pulsed supersonic beam of cold YbF. This is far superior to the previous effusive oven source (see Hudson [16]); not only is a higher flux of ground state molecules observed, the running of the source is far simpler³. The development of the source has been extensively documented by Condylis [14].

Ytterbium fluoride (YbF) - a radical - is very reactive, so to make a YbF beam the molecules must be made on demand. To do this precursors (Yb Metal and SF₆) need to be heated strongly

³The reader should consider that in making an improved measurement we are concerned with improved sensitivity per root time. Whilst some developments are designed to improve signal to noise, advances can also be made also by simplifying running procedure. For example, the pulsed supersonic source typically requires only a few minutes attention per day, compared with several hours using the oven source.

using a pulsed Nd-YAG laser. However, once formed, the molecules need to be cooled so that there is a significant ground state population. To solve these conflicting constraints laser ablation is combined with supersonic expansion cooling.

A dilute mixture of approximately 2% SF_6 to 98% Ar is introduced into the chamber in very short pulses ($\sim 50 - 100 \mu\text{s}$) using a commercial pulsed solenoid valve⁴. The pressure in the source valve changes from 4 bar to 10^{-6} bar in a few μs resulting in rapid cooling of the gas pulse. The argon carrier gas dictates the final thermal properties of the pulse, whilst the SF_6 provides the fluorine for the YbF molecule. We find higher SF_6 concentration leading to reduced YbF flux. Using a pulsed Nd-YAG laser, a carefully timed laser pulse is used to ablate a small amount of ytterbium metal into the cold gas pulse. This reacts readily with the SF_6 to form YbF. The hot YbF then thermalises with the carrier gas to produce a cold pulse of YbF, analysis of which suggests a temperature of approximately 5 K and a velocity of 530 ms^{-1} . The translational temperature is calculated by measuring the YbF pulse width at the end of the machine using the PMT. Under the assumption that all molecules in the pulse were created simultaneously at a single point in space, we can calculate the temperature from the distribution of observed velocities. This was fully investigated by Condylis[14], after implementation of this source. The molecular pulse then passes through a skimmer into the upper vacuum chamber. The rôle of a skimmer is to allow the central, most intense part of the YbF pulse to pass into the upper chamber, whilst the vast majority of the $\text{SF}_6 + \text{Ar}$ remains in the lower chamber, to be pumped by two large turbo pumps. This results in a clean molecular pulse which can be inspected either using LIF (probing the YbF using a resonant laser), or using a fast ion gauge (probing the entire gas packet by measuring minute changes in the vacuum as the pulse passes).

Other carrier gases including the heavier xenon have been tried. This makes the beam slower (approximately 50%), affording a longer interferometer coherence time. However, the transverse velocity distribution of the pulse remains constant, and leads to an increase in solid angle subtended by the beam. Taking EDM data with Xenon carrier gas resulted in no net improvement in sensitivity when acquiring EDM data.

The ytterbium target is formed from a ribbon of ytterbium glued with *torr-seal*^{footnote}Torr-seal is a vacuum compatible adhesive. to a steel disk approximately 12 cm in diameter. The surface condition of the ytterbium metal strongly influences the quality of the molecular beam, so periodically the target is rotated to expose untarnished ytterbium. The lifetime of the target varies considerably, but typically we see that after a few weeks of intense running it is very hard to obtain a satisfactorily quiet YbF beam. At this point the target is removed and resurfaced on a lathe. This restores the source to its original condition. After 3-4 of iterations insufficient YbF remains, so after cleaning the steel disk, a new target can be fabricated by attaching fresh ytterbium.

⁴It is worth noting that though we use a commercial valve, we operate well beyond the usual operating regime. Instead of applying 5-10 V to open the valve ‘slowly’ we pulse 300-400 V for a period of approximately 200 μs

The quality of the molecular beam is strongly dependant on the condition of the ytterbium, and as a result there is no formulaic method of source optimisation. However, the timings of the pulsed gas valve, the ablation laser and the ablation laser power (controlled by tuning the delay between the flash lamps and the Q-switch) are largely orthogonal. Adjusting each source parameter to give the most intense beam usually has fair results (see appendix A.1 for optimal timings). The correct measure for the *optimal* molecular beam is not purely intensity or stability⁵, but a combination of sensitivity and longevity. Whilst running we aim to optimise sensitivity per root day. This typically means maximising the sensitivity over contiguous data units (clusters) of about 30 minutes. Of course this measure contains not only molecular beam noise, but also magnetic field noise, low PMT noise etc. Despite this, in practice this is the most useful measure. Note that it is often seen that the most intense beam will decay into a noisy beam more rapidly than a small initial beam, so some experience is required for efficient operation. In section 3.7.2 the effect of an unstable source is discussed more fully.

2.1 Lasers

The various laser systems used have changed little since the previous experiments by Condylis [14]. In the most basic configuration we require two lasers: the Nd-YAG for ablation and a source of 553 nm continuous wave light for pump/probe purposes.

2.1.1 Ablation laser

The ablation laser is a Brilliant Nd-YAG modified to run at 50 Hz. In actual fact it is typically run at 25 Hz as although the higher repetition rate should double the data acquisition rate, no net improvement in sensitivity was observed, due to insufficient vacuum pumping⁶. Since the purpose of this laser is purely to heat the surface of an ytterbium target, creating a hot vapour plume, the laser wavelength is largely unimportant. The only parameter that is routinely optimised is the pulse power, controlled by the delay time between the flash pump and the Q-switch. All experiment timings are reference from the Nd-YAG Q-switch.

The focus and positioning of the ablation laser has a large effect on beam quality. The positioning of the Nd-YAG laser beam upon the ytterbium target in relation to the gas valve must be correctly set. The ablation beam aligned above the pulsed valve aperture, with the beam centre aligned horizontally on the Yb disk⁷. The focus of the ablation laser may also be modified using a simple telescope. We find a slight reduction of the spot size provides the optimal molecular beam.

⁵Which one might choose to calculate as the variance of integrated intensity over a series of shots, or perhaps some long term drift

⁶Insufficient vacuum pumping, resulting in a higher source chamber pressure tends to severely attenuate the beam.

⁷Very intense (at least a factor of two gain in observed flux), yet unstable molecular beams can be made by positioning the ablation laser 2-3 cm to the side of the pulsed gas valve.

The final spot is approximately 5 mm in diameter at the target.

From day to day the optimal parameters tend to drift as the target condition varies. Typical values for source parameters are given in appendix A.1, whilst schematics describing the source assembly can be found in Condylis’ thesis [14]. Detailed CAD files may be found in the Centre for Cold Matter CAD repository.

2.1.2 Pump - probe laser

The 553 nm pump - probe is most demanding laser system. A single mode dye laser is used to generate the required light. As with Hudson and Condylis’ work, a Spectra Physics 2580 Ar⁺ laser is used to pump a Spectra Physics 380D ring cavity dye laser. For this wavelength, Rhodamine 110 dye is used, pumped with 3-5 W Ar⁺ light, generating⁸ 200-300 mW at 553 nm. Typically we achieve dye lifetimes of 2-3 weeks with moderate use, falling to less than 1 week during intensive data taking.

Two stages of lock are used. The cavity of the 380D is stabilised using a commercial Spectra Physics reference station, which contains of two temperature stabilised Fabry-Perot cavities. This stabilises the dye laser to 500 kHz on times scales of order 1 second, and allows smooth scans over several GHz. However, the stabilisation cavities have a slow drift of 5 MHz/h; for EDM data acquisition we need long term stability of order a few megahertz over 12 hours, less than the linewidth of the Q(0)F=1 pump/probe transition. To provide this the cavities are locked to an I₂ spectrometer (see Hudson [16] 2.4, p. 31).

The I₂ spectrometer implements saturation spectroscopy of iodine. Overlapping counter-propagating pump and probe beams are passed through an iodine cell. The pump beam is frequency shifted using an acousto-optic modulator (AOM) by a tuning frequency 254 MHz, overlapping the optical transition in Yb with a strong I₂ saturation feature. Once optimised and thermally stable, the laser will remain locked for up to 12 hours⁹.

2.2 Electric fields

To measure the electron EDM, a strong electric field is required to polarise the molecules. To characterise the field we consider:

- absolute value of field
- quality of field reversal
- spatial homogeneity of field over a large region

⁸This level of output is required to reliably lock to the iodine spectrometer.

⁹Though quite robust, a busy lab or amplitude modulations on the laser can lead to the laser unlocking more frequently, such that it is often the limiting factor when taking long (30 minute) data runs.

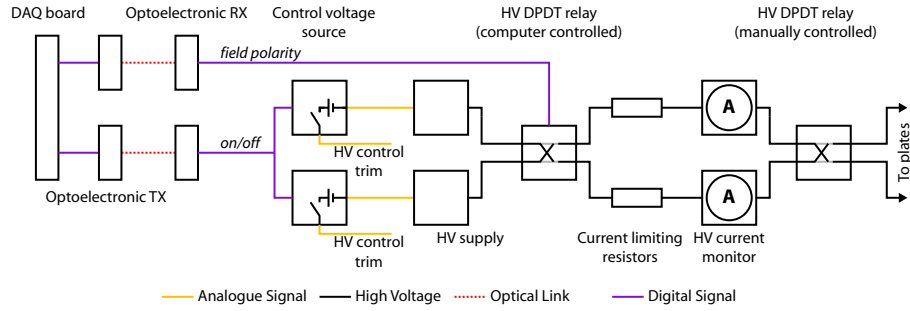


Figure 2.2: Schematic of electric field power supply control. The HV control system allows the electric field to be switched on/off remotely, as well as reversing the field polarity.

- leakage current

It is reasonably straightforward to generate very high inhomogeneous fields, but it is hard to control homogeneity and leakage current. High voltages are also hard to switch fast without causing large charging currents.

The electric field systems are most naturally considered as two sections: HV power supplies and electric field plates.

2.2.1 Power supplies and control

We use two 50 kV Bertan power supplies for the EDM experiment. The control requirements are simply to be able to turn off the electric field, and to reverse the field polarity. We obviously need to be able to tune the potentials applied to the plates, as for a well supported EDM measurement we will need to take data at various voltages, to ensure the measured EDM phase scales with the molecular polarisation η . However, on timescales of a day, we do not need to dither the voltage. The power supply voltage is governed by a 0-5 V DC control signal. At present we derive this from a 5 V reference and a potential divider. The output voltage of the supplies can be set to zero using a relay to drop the control voltage to 0 V. We must do this every time we reverse the electric field direction. The polarity of the electric field is reversed using large pneumatic HV relays. Both the polarity and power supply state (on/off) are remotely controlled by the computer. These control lines are electrically isolated, using optical control over multimode fibre, to prevent systematics arising from control line currents, as well as isolating the computer from HV spikes that may occur. Large series resistors are placed on the output of the HV relays. These are chosen to reduce charging currents to a safe level (to prevent magnetisation of the experiment, dependant on the state of the electric field), though obviously at the expense of increased electric field switching time. These systems are shown in figure 2.2.

A second manual ‘relay’ is used to periodically ($\sim 1/\text{hour}$) reverse the polarity of the field. This reversal is respected by the analysis, but is not controlled by the computer. It reverses relationship between the logical state of the computer generated electric field control signal and the actual

direction of the field. This reversal differentiates between two types of measured ‘EDM’. A genuine EDM signal, will necessarily change sign according to the state of this switch. A fake EDM - say from a magnetic field created by the HV supplies - might appear similar to an EDM, but would not be correctly correlated with the manual reversals.

2.2.2 High precision field plates

The most obvious constraint is that the electric field must be sufficiently uniform that the entire molecular pulse is on resonance. Given a FWHM linewidth¹⁰ of 44 kHz at 18 μ s pulse length, electric field of 20 kVcm⁻¹ and plate spacings of nominally 1.2 cm, the required engineering tolerances are $\pm 115 \mu$ m over each pulse region.

The electric field between the rf transitions is less of a concern, as the interferometer averages over the integrated electric and magnetic fields between rf pulses. In the work of Condylis [14] three regions of electric field were used, with transitions occurring in so called ‘guard’ (g) regions, and the majority of phase being acquired in a higher voltage ‘centre’ (c) region. However, this scheme has since been simplified so that only one region is used, with the entire experiment occurring in a single field region. This new scheme is designed to suppress certain systematic effects which can occur as the molecules rotate between regions of different electric field (see section 4.2.1).

To form a single region of homogeneous field, two parallel aluminium plates are used. The physical dimensions of the plates were constrained somewhat by the machine dimensions, maximised to fit into the upper vacuum assembly, whilst the width is such that we can be sure that edge effects may be neglected. The plate separation was chosen to allow room for the molecular beam to propagate, given that the pulse has some small initial thermal velocity distribution¹¹. The final plates were 750 mm long, separated by 12 mm, yielding a useful interaction region of at least 700 mm.

The construction of the field plates was technically demanding. At the required tolerances, very careful techniques were needed. The plates and plate supports were constructed from aluminium tooling plate, which resists deformation when machined. All alignment holes were made using a wire eroder¹². The plates were electro-polished to remove small surface features, before being gold plated¹³. The plates are shown in figure 2.3 (gold coloured), along with the support structure. The end dowels (white) are made from ground alumina. These precisely align the plates. The remaining green dowels are lower precision torlon. These primarily add strength to the structure. As well as a CAD rendering (figure 2.3), photographs taken during the plate upgrade process are shown in appendix E.

Once installed inside the machine the plates were conditioned to support higher electric fields.

¹⁰The rf transition is discussed more fully in section 3.4.

¹¹As an rough guide, the field may be considered uniform once the molecules are several plate spacings away from the edges. The assembly was modelled in 3D in FEMLAB to allow us to derive an accurate model of the field.

¹²Wire eroding is used to bore high precision holes through metal without any mechanical stress.

¹³Care was taken to gold plate without the commonly used ferromagnetic nickel layer.




Figure 2.3: High voltage plate assembly. This recent rendering is of the single region design. The assembly shown here is mounted vertically inside the chamber, the molecular beam propagating freely between the plates along the major axis.

The voltage was increased until a small (20 nA) current was detected using the analogue leakage current monitors. A script was used to repeatedly switch field polarity every minute or so. Over time, the leakage current was seen to die away, at which point the field was increased. The conditioning procedure can take several weeks. The end result was that the operating field was increased from 13 kVcm^{-1} , to 20 kVcm^{-1} , with less than 1 nA of leakage current.

To connect the plates to the power supplies, HV feed-throughs are mounted on the side of the vacuum chamber. Spring loaded probes connect these to the plates under vacuum. The square slot cut into the support structure to connect the HV to the plates can be seen in figure 2.3.

Throughout the rest of this thesis I will refer to the section of molecular beam which is within the region of homogeneous electric field as the ‘interaction region’. It is here that we drive rf transitions, and crucially, where the molecules are during the free evolution phase of the interferometer.

2.3 Radiofrequency fields

The rf fields have been extensively developed throughout the course of my research, moving from loop antennae to a transmission line configuration. The EDM data presented in this thesis was acquired using the rf transmission line, though rf loops are discussed in the systematics chapter.

The rf pulses used to drive the π -pulses for the interferometer have four main parameters: amplitude, duration, timing and frequency. These are all controllable through careful synthesis and amplification. Our biggest worry is that poor control of the rf transitions may lead to a systematic effect or contribute to the noise in our measurement. For the time being it is sufficient

to bear in mind that we wish to have a handle on these parameters, and to note where we fail to obtain full control.

2.3.1 Radiofrequency structures

Radiofrequency loops

The conventional configuration for creating an rf field to drive a magnetic dipole transition an rf loop antenna. This approach can be traced back to Ramsey’s original neutron interferometer experiment [17]. The distribution of the rf field created is that of a DC magnetic field from a single coil. This was used in the work of Hudson [16], as well as in the work of other groups. The machine was equipped with 4 loops, of 8 cm diameter, placed along the beam line, though only the loops on the end of the interaction region were used. Our original scheme, implemented by Hudson, utilised long rf pulses, such that the distance a molecule moved during the pulse was comparable to the extent of the a short 5 cm region of homogeneous electric field. The rf pulse length was fixed the length of time the molecules remained on resonance, as they passed though the brief guard region electric field. The π -pulse was tuned using the rf synthesiser output power. In this regime, all molecules across the pulse can integrate a similar field, and receive an identical transition.

Our current experiment uses short rf pulses, which reduce systematic effects associated with propagation through spatially varying electric fields. In the short pulse limit, the molecules travel a much shorter distance during the rf pulse. In particular we can use a single region electric field design using this scheme. However, the downside is that when using loops with short rf pulses, molecules at different positions in the YbF pulse will see very different rf powers, leading to an inhomogeneous rf transitions.

In practice we were able to drive π -pulse transitions with 40 dBm rf power (10 W), in 18 μ s (over which time the molecules move ~ 1 cm).

Transmission line rf

More recently an improved rf structure has been implemented. Using short rf pulses with rf loops we cannot drive homogeneous π -pulse transitions across the molecular pulse. Loops also tie the rf transitions to one point in space - it is not possible to drive transitions further into the interaction region for example. Our novel solution was to use the electric field plates as an rf transmission line. The plates in cross-section appear as two cores, each 70×8 mm, separated by 12 mm. The parallel plates act as a transmission line, supporting a transverse mode (TEM) rf field. In implementing such a structure the rf needs to propagate from the amplifier, through the vacuum assembly and shielding, into the transmission line, and out again into a 20 W terminator. The impedance of this transmission line was calculated by measuring reflected power, and found to be 35.5 Ω . The rf was




Figure 2.4: Transmission line rf antenna and coupling capacitors. The HV plates are connected to the external rf systems through HV capacitors, whilst trimming capacitors are used to match the impedance of the plates to the rf line.

fed from the final power amplifier to the machine using low loss $50\ \Omega$ coaxial cable, though the vacuum assembly, into non-magnetic semi rigid¹⁴ (RG-405). To couple to the plates a capacitor network was used to isolate from the HV (see figure 2.4.). Trimming capacitors were added to impedance match the $35.5\ \Omega$ transmission line to the $50\ \Omega$ semi-rigid coaxial cable, minimising reflections¹⁵. The transmission line rf plates reflect 20% of the input power due to the remaining impedance mismatch. This leads to a standing wave throughout the machine, though the rf field is still much more uniform than when using loops. This is discussed further in section 5.1.4, where a comparison of field homogeneity is made in figure 5.11.

2.3.2 Radiofrequency mixers, synthesiser and amplifier

To drive appropriate rf transitions we need good control of the timing, length, amplitude and frequency of the rf pulse. These requirements are independent of whether we use loops or the transmission line implementation. However, in each case the rf network is slightly different because the transmission line system has one rf field, pulsed twice, whereas using two loops, two distinct rf fields are created.

Radio frequency equipment for use with loops

Under ordinary running conditions we use two rf transitions which split and recombine our interferometer ‘arms’¹⁶. The rf source is an HP8657A synthesiser, programmable over GPIB. Typically, the GPIB reprogram time is of the order 100 ms. Since the molecules fly between the loops in roughly 1 ms, this is too slow to reprogram between the upper and lower loop, so another strategy is needed to tune amplitude, gate the pulse and select the frequency.

This solution is shown in figure 2.5. The output from the synthesiser is immediately split into two channels. Each channel now passes through a switch (to gate the CW rf to a pulse), then

¹⁴Care must be taken to use non-magnetic semi rigid, as commonly the construction is to use a silver plated steel core.

¹⁵The $35.5\ \Omega$ impedance is due to the plate geometry. In future plates should be designed with the geometry tuned such the resulting impedance is $50\ \Omega$.

¹⁶No more than two loops are ever used, the extra loops are for investigating different spatial regions inside the machine.

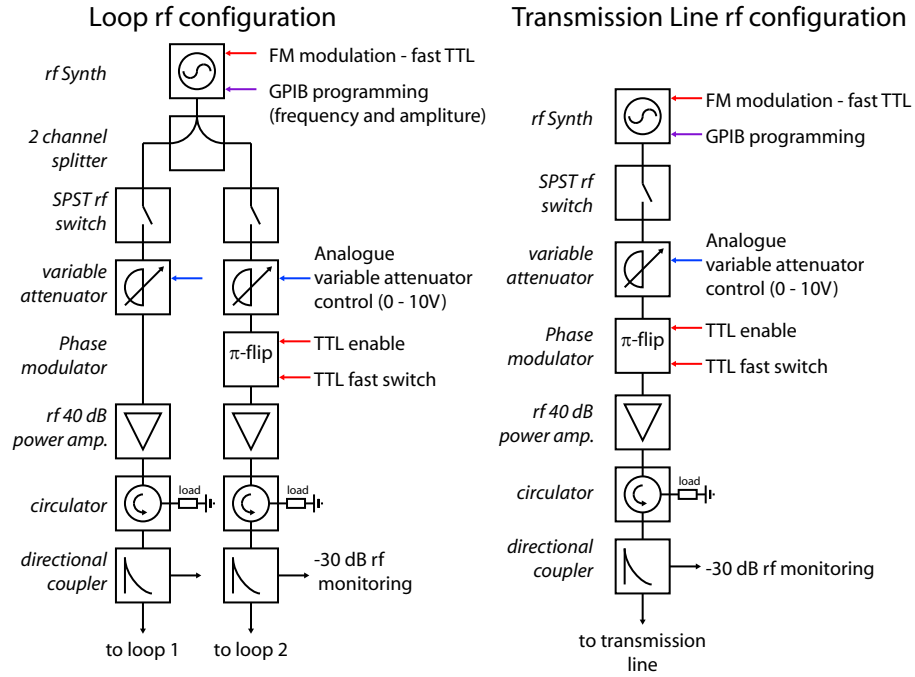


Figure 2.5: The rf configuration used in measuring the EDM. Slightly different layouts are required, depending on whether loops or the transmission line are used to generate the required rf.

a attenuator (to independently tune rf power). The switch is controlled by a fast digital line, giving μs timing resolution. The rf amplitude is adjusted using a manually controlled variable attenuator. The second channel has a phase flipper, which can switch the phase of the rf between 0 and π . This again is controlled by a fast digital line¹⁷. Each signal now passes through a 40 dB amplifier, before -30 dB is picked off for monitoring using a directional coupler. A circulator is used to protect the amplifiers from reflections.

This configuration gives us control of the amplitude of the top and bottom rf field and the start and end time of each pulse (to the μs level). Using the synthesiser frequency modulation (fm), we can use a digital control signal to toggle a frequency offset, fast enough to tune the top and bottom pulses independently¹⁸. Beyond this the amplitude and frequency are fixed on the shot timescale. Although GPIB limits the repetition rate to 10 Hz, we can, running our source at lower frequency, increment parameters over successive shots to explore the rf parameter space - see section 3.2.

Radio frequency equipment for use with the transmission line

The transmission line rf equipment differs slightly from the loop configuration in that there is only one rf field. The most straightforward configuration is shown in figure 2.5. It consists of a single synthesiser, digitally controlled fm, variable attenuator and digitally controlled switch. Again, the switch gates the CW synthesiser output into two short pulses, whose timings are controlled by

¹⁷Periodically shifting the phase is used later as a probe for systematics involving incomplete pumping, and the $|F=0\rangle \leftrightarrow |F=1\rangle$ coherence (§3.6.3.)

¹⁸Although the electric field plates are designed to be parallel, engineering tolerances limit the precision to around $100\text{ }\mu\text{m}$. This means that the relative rf transition frequencies may be up to 23 kHz detuned when running at fields of 20 kVcm^{-1} . This is comparable to the linewidth of the rf transition.

the pattern generator. It is also possible to program the synthesiser fm, which can then be toggled using a fast digital input. This allows independent control of the frequency of each pulse. This configuration does not allow for independent amplitude control of each pulse¹⁹. In practice this is inconsequential as typically rf amplitudes for each transition are very similar. As before we can shift the phase of the rf by π between pulses if needed.

2.4 Static magnetic field

To measure the EDM robustly we require very precise control over the static magnetic field, particularly in the direction to the static electric field.

2.4.1 Magnetic field coils

As shown in figure 1.10, we measure the EDM by comparing two interference scans (scanning B) taken at different electric fields.

From chapter 1, the interferometer lineshape is $R_0 \sin^2(\phi_{de} + \phi_B)$. Fixing the EDM phase, we can move the interferometer output through the application of a small magnetic field. Recalling that $\phi_B = 2\mu_B Bt/\hbar$, and an interferometer interaction time of 1000 μ s, to shift the interferometer output 1 fringe (π) corresponds to a field of 34 nT.

The field coil is formed from 4 wires glued to the side of the inner shield, in the configuration shown in figure 2.6. This is unchanged since implementation by Redgrave ([44], p. 37.). At this time, despite the close proximity of the μ -metal shielding, it was shown to give rise to a field uniform to one percent throughout the interaction region²⁰. We generate a field of 16 nT/mA, so to scan a few fringes takes approximately 10 mA. The exact calibration of the interferometer to jump from one fringe to the next is carried out empirically on a regular basis.

The nature of the experiment necessitates very rigorous control. The biggest risk is that magnetic field through which the molecules pass changes with the state of the applied electric field. This will be discussed further, after the introduction of a data acquisition schemes (chapter 3).

Two current sources are used. The most simple converts a voltage (0 to 5 V) to a current, typically ± 5 mA. Using this we can scan a range of B fields by varying the output voltage (derived from a National Instruments analogue output). This is particularly useful when performing interference scans, producing fringes similar to those shown in figure 1.10. Interference scans are discussed in section 3.6.

For taking EDM data this current source is inappropriate, as reliable currents would be hard

¹⁹In practice we can still tune the rf pulse length such that the integrated power for each pulse may be independently tuned, though this approach introduces unpleasant asymmetries in how we deal with the split / combine stages of the interferometer.

²⁰In chapter 5.1 I discuss methods recently developed to map magnetic fields inside the experiment.




Figure 2.6: The magnetic field coils consist of two single turn loops glued to the inside of the μ -metal shielding. These generate a field of 16 nT/mA.

to establish. In practice we sample a small set of points from the interference curve, at 4 magnetic fields. For this purpose, a current stepping supply is used. This takes two balanced optical digital inputs B and ΔB , corresponding to steps in current, the magnitude of which can be tuned using potentiometers. The linear combination of these steps with a constant offset gives the four magnetic field states required for EDM data acquisition. The data acquisition scheme is developed fully in the following chapter, section 3.6.1.

Optical isolation is used to carry signals across the lab without risk of generating magnetic fields. The implementation is straightforward: the digital output from the computer / instrument is used to drive an optoelectronic driver. Broadband plastic fibre is used to carry the light. A standard optoelectronic receiver unit then converts back to an electrical signal. In some circumstances (e.g. electric field state control) a parallel, inverted copy of the digital line is created, such that the net field from the Tx/Rx units is zero. Optical isolation is also used for electric field control.

As with the electric field we perform manual magnetic field reversals, reversing the relation between the computer field state and the direction of the electric field. These have the effect of shifting the interference curve half a fringe ($\pi/2$), changing the sign of the gradient of the fringes, and hence that of any measured EDM. Again, our final EDM analysis must be calculated to reflect this.

2.4.2 Magnetic shielding

The Zeeman interaction is many orders of magnitude stronger than the EDM interaction. Magnetic fields affect the experiment in several ways. Firstly Zeeman splitting can broaden the $|F = 0\rangle \longrightarrow |F = 1\rangle$ transition, leading to partial π - pulses. Given the linewidth of the transition, approximately 44 kHz, the ambient internal magnetic field must be less than 0.5 μ T to suppress broadening. Given a typical lab field of 50 μ T, a shielding factor of at least 100 is required. Secondly, random magnetic

field noise will limit our sensitivity [44]. Improved shielding will help reject magnetic field noise.

Two layers of concentric μ -metal shielding are used. The outer shield (64 cm \times 110 cm) is a large open cylinder, which encases the entire vacuum assembly, detectors and turbo pumps. The inner shield (17 cm \times 79 cm) is positioned inside the vacuum chamber, encasing the rf loops and field plates. The entire interaction region lies within this shield.

The overall effect of the shields is to reduce the z component of the magnetic field (B_z) nominally by a factor of 1000 throughout the interaction region. Ambient fields in the interaction region are typically less than 5 nT. The vertical component B_y is shielded poorly, due to the location of the access holes, and the absence of end caps on the outer shield. This component is much less important however, except in the consideration of certain systematic effects where rotation of the molecules renders us sensitive to magnetic fields other than B_z .

2.5 Instrumentation

2.5.1 Photomultiplier tubes and laser induced fluorescence

Our primary detection mechanism is laser induced fluorescence (LIF), a direct probe of the population of a given state.

Our EDM experiment demands that after the interferometer we can measure the population of either the ground state $F = 0$, or $F = 1$ levels, where the interferometer phase can be implied from either measurement (equation 1.3.e). The LIF scheme has changed little since the early YbF spectroscopy by Redgrave [44], and later Hudson [16]. We probe $F=1$, on the $Q(0)F=1$ transition. A probe laser drives the transition $X^2\Sigma^+(\nu = 0, N = 0) F = 1 \longrightarrow A^2\Pi_{\frac{1}{2}}(\nu = 0, N = 0)$. A photomultiplier tube (PMT) is used to measure fluorescence as the molecules spontaneously decay from this state. The PMT signal is then proportional to the population of the $F=1$ state.

We also have a second ‘normalisation’ PMT which detects fluorescence at the pump stage of the interferometer. This allows us to normalise our interferometer output to the beam intensity, removing some of the source noise. This scheme has yet to be fully implemented, and the EDM data presented here is not normalised.

For certain tasks such as source optimisation the pump beam (discussed in §1.3) is not used, so the probe PMT measures the population of the $|1, \pm 1\rangle$ states unpumped. This population is proportional to the ground state $|0\rangle$ population, so we can tune the source parameters independently of the rf parameters. For rf tuning the pump block is removed. The rf repopulates the $|1, \pm 1\rangle$ states, which is detected by the probe. This constitutes a single resonance rf experiment, and is discussed fully in section 3.4. Observation of the LIF signal in this configuration allows the rf parameters to be tuned.

2.5.2 Magnetometry

Using our interferometry scheme, a single EDM measurement takes at least 20 seconds to record (dominated by the switching time of the electric field), though in fact our more elaborate scheme takes nearly 3 minutes. During this time the magnetic field can vary significantly, introducing noise in the measurement (discussed later in section 3.7.3). A Bartington 3-axis flux gate magnetometer (Mag-03MCL100) is used to monitor these ambient magnetic field changes. This is placed between the magnetic shields, outside the vacuum chamber (it is not presently practical to place a magnetometer inside the vacuum). It gives us real time magnetometry with nano-Tesla sensitivity between the shields. This of course has the limitation that although sensitive, we are not monitoring exactly where the molecules are. Field gradients, caused by nearby equipment, certainly add some inhomogeneity to the lab field vector. Section 5.1.3 presents a method to map static internal magnetic fields.

The rôle of the magnetometer is not to correct for Zeeman induced interferometer noise. As we are sampling the incorrect region of space, such an approach would not be robust. Instead we use the magnetic field information to allow us to selectively discard data taken when the magnetic field was most noisy. The veto of noisy blocks is discussed in section 3.7.3.

Leakage Monitors

As part of controlling magnetic fields in the interaction region, we have to be careful to control the current flow when charging electric field plates. Our strategy is to allow charging currents to decay to negligible levels before data acquisition in a given electric field polarity commences.

Current monitors on the HV cables allow leakage and charging currents to be recorded. These simply monitor the potential difference across a large $10\text{ M}\Omega$ sense resistor. We can measure leakage currents to a precision of approximately nano Amp. The current leakage monitors display information on an analogue panel metre; this information is not yet recorded by the computer. In running the experiment we occasionally monitor currents to ensure they are below ‘safe’ levels.

A safe level is typically argued to be 2 nA. This is based upon Hudson’s reasoning (see [16].3.6.1) that 10 nA could in the worst case give rise to systematic EDMs of order $4 \times 10^{-28}\text{ e.cm}$; we suppress current leakage systematic EDMs to an order of magnitude below Commins’ experimental limit of $8 \times 10^{-28}\text{ e.cm}$

2.5.3 Data acquisition

An experiment of this complexity needs a high level of computer control. The software developed is complex, so full listings of the code are omitted. What follows is a basic summary of our approach, detailing our choices of hardware, and any implications these may have for the experiment.

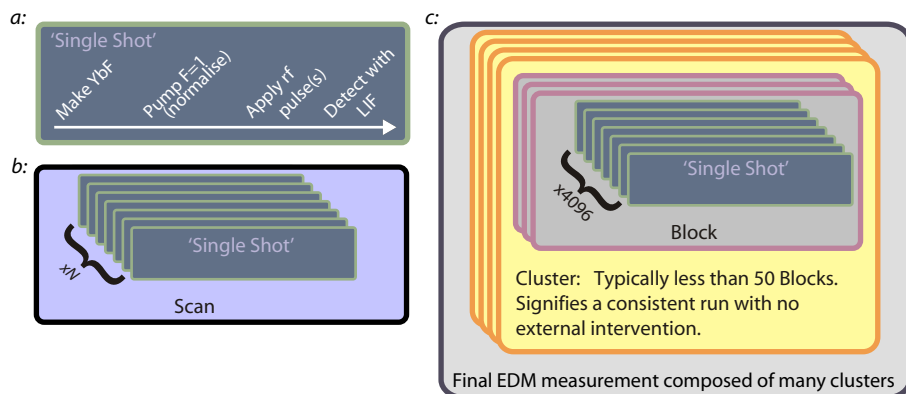


Figure 2.7: Data structures associated with the EDM experiment: *a*: The most simple data structure, a shot, contains a single TOF, although several operations may have been performed on the molecules. *b*: A series of shots are compiled into a scan to explore a parameter space. *c*: The EDM data acquisition scheme takes a series of shots at various electric and magnetic fields. These form a block. Several blocks form a cluster. The set of all clusters make our final EDM measurement.

Data structure

So far I have introduced the interferometer lineshape (§1.10), and discussed the implementation of a pulsed source of YbF (§2.0.2). Let us now consider the data structures required to build up the basic.

Using a pulsed source, the most fundamental task we can perform is to make a single YbF pulse. This is termed a *shot*. Whatever we do to the YbF (ranging from nothing (perhaps to align the probe laser) to a multiple rf pulse experiment, we fundamentally measure one thing about it - the time of flight (TOF) using LIF and the probe PMT. The PMT is followed by a fast current to voltage preamp.

Taking a series of shots we can adjust one or more aspects of the experiment to explore parameter space. These might be source parameters, rf parameters or field parameters. We call this element a *scan*. Whilst taking EDM data we do something very similar to this, but in order to optimise our use of the lineshape we scan discrete set of points²¹. This special type of EDM data acquisition scan is called a *block*. Each block consists of 4096 individual shots. We consider the block the smallest data unit which contains a reliable EDM measurement, because of the order in which we scan points. However, of the 4096 shots taken, there is some replication, so in fact each block contains 16 sub-blocks - each containing an EDM measurement. This is discussed fully in chapter 3. To make a more robust EDM measurement we will have to aggregate blocks into *clusters*, and clusters into a final EDM data set. Assuming that data acquisition is running well, clusters are typically 20-50 blocks long. We take the opportunity between clusters to check magnetic field bias, and perform manual field reversals. Nestings of these structures can be seen in figure 2.7.

²¹This is fully discussed in chapter 3.

Data acquisition and experiment control

Our experiment is controlled by a single computer, with data acquisition carried out using a selection of National Instrument (NI) boards and USB peripherals. Perhaps the most obvious way to address such hardware is using NI Labview. However, use Measurement Studio, a set of libraries designed to control every aspect of NI's data acquisition hardware, similarly to Labview. Measurement studio libraries are called directly from C#, Microsoft's object orientated language developed as part of the .NET environment. Once familiar with the concepts of the language this affords more flexibility than Labview, is more robust, and is certainly more appropriate for developing a complicated acquisition scheme. The compromise is that the development of these systems requires an advanced level of programming skill.

The EDM computer is equipped with a range of PCI and USB devices NI-DAQ devices. A block diagram of the hardware architecture is shown in figure 2.8. The constraints of the experiment define some of the hardware/software implementation. The timescale of a single shot is 1 ms, with events occurring on the 1 μ s scale. This necessitates hardware controlled timing. Anything that happens on the shot timescale must be triggered using the pattern generator, which can be programmed with microsecond precision. For example the pattern generator creates a timing edge for the Nd-YAG Q-switch, all valve timings and rf modulations. A key advantage of the pattern generators are that even with precise timing and triggering requirement, they add very little computational load. This is very important, as if the EDM data acquisition were to fail due to insufficient resources in a systematic manner, then we would likely 'measure' a systematic EDM. The pattern generator also provides the main reference clock for the experiment, which itself is software phase locked to the mains frequency.

The inter-shot time scale is several orders of magnitude larger (40000 μ s). Here we can reprogram²² the rf synthesiser using GPIB²³. This allows us to record a series of shots (scan) at different rf frequencies, tracing out an rf transition. We use RS-232 serial communications to program the Brilliant Nd-YAG.

Key to data acquisition is the sampling of the PMT and magnetometer. The sampled PMT signal is recorded in μ s bins. Gates are set so that we only record when we expect the pulse to arrive at the probe region. In ordinary scans no analysis is performed in realtime. We can later load²⁴ the blocks into *Mathematica*, retrieving information on all the instrumentation detailed here, as well as experimental parameters. The EDM data acquisition mode is a little more focused. It provide realtime analysis of the EDM data at a block level. This is essential; it indicates that the interferometer is correctly configured, the laser is locked and if the experiment is excessively noisy.

²²The realtime reprogramming of the rf synthesiser is quite slow, and necessitates the experiment running at 10 Hz. Even then, there are occasional rf drop-outs. This is acceptable for scanning the rf lineshape, but inappropriate for taking EDM data.

²³GPIB is a hardware control protocol used to link test and measurement instruments.

²⁴More formally, the scans are serialised .NET objects which must be properly decoded using the EDM control software and *Mathematica* .NETLink.

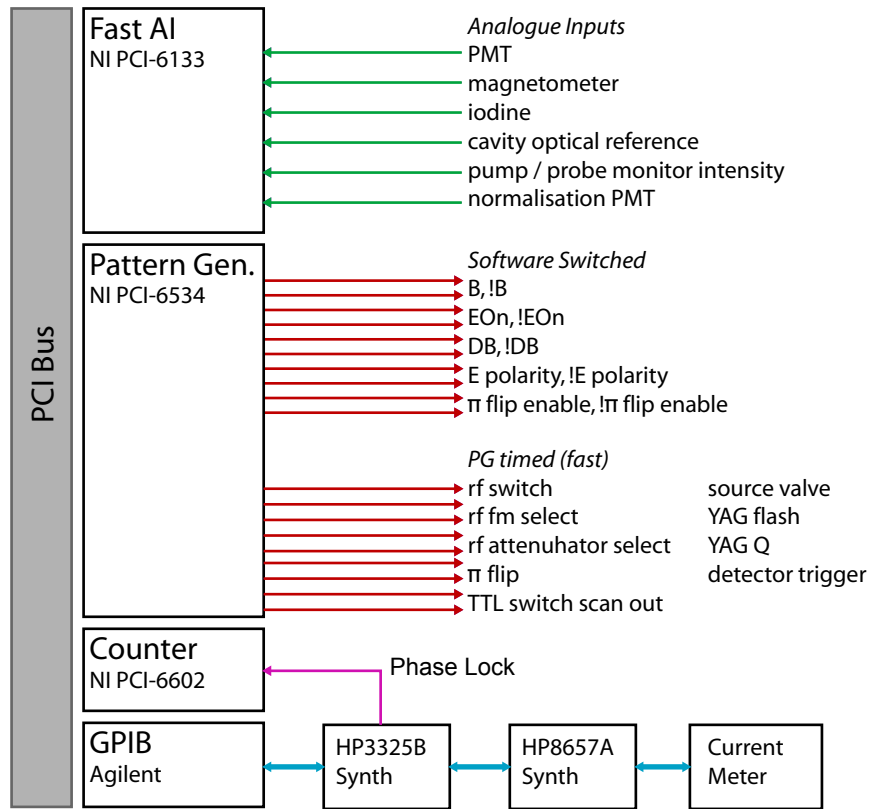


Figure 2.8: Computer hardware I/O.

The blocks are later analysed more thoroughly. Typically, damaged blocks will be rejected, before a database is compiled in *Mathematica*. This can then be analysed in order to yield the complete EDM analysis. This process is detailed in chapter 6.

Data acquired using the analogue inputs has units of volts, so TOF data is a time varying voltage. Integrated TOF signals have units of $V\mu s$.

Chapter 3

Experimental Method

Chapters 1 and 2 outline the basic principle of how YbF is used to measure the electron EDM and give a brief overview of the apparatus used. A more detailed experimental method is now presented. Before discussion of the EDM experiment, I explain the operation of the beam and interferometer.

3.1 The molecular beam

Fundamental to the experiment is the creation of a stable molecular beam. This brief section will tie together many of the technologies discussed previously.

The first step is to stabilise and lock the probe laser so that we can detect the molecules with laser induced fluorescence (LIF)¹. The dye laser is first slaved to a stable cavity [16]. Assuming the laser cavity and stable cavities are well aligned this simply involves throwing the stable-lock switch on the dye laser control. To lock to the iodine spectrometer the laser frequency is scanned. Monitoring transmitted probe light through the iodine cell we see a series of distinct dispersion features. We then fine tune the laser frequency and, by eye, lock to a specific feature.

As introduced in the previous chapter, the most elementary unit of data considered is a *shot*, which consists of flashing the Nd-YAG pump, opening the gas valve, firing the Q-switch, then some time later, acquiring the time of flight² (TOF) signal using laser induced fluorescence (LIF). The most simple profile³ would simply repeat this sequence, recording a series of TOF profiles. All other tasks, with the exception of actually acquiring EDM data, are generalisations of this process.

We use this scheme to align the probe laser. A series of shots are generated, with all controllable parameters held constant. The probe alignment is then adjusted whilst monitoring the integrated

¹Throughout this chapter I will not be discussing mundane aspects of operation, switching equipment on and the like - a detailed operating procedure is however included in appendix D.

²The TOF is discussed in more detail in the next section.

³In order to perform different types of scan, the control software contains a set of configurable tasks, called *profiles*. These are generally for scanning some parameter space. They contain information on how the experiment is configured, as well as which parameter should be scanned.

TOF. Monitoring this, we maximise the peak intensity, whilst being careful not to generate a large background of scattered probe light.

The quality of the molecular beam is best assessed by acquiring EDM data and comparing the EDM sensitivity to the shot noise limit for a given sample size⁴. In practice, to start the experiment we usually spend some time at the beginning of each day optimising the source by firing a series of shots as discussed above. We generally look for a large beam, with a clean gaussian velocity distribution. Looking at the integrated time of flight series, it is possible to see qualitatively if the beam is very noisy, in which case the target is moved to expose ‘fresh’ ytterbium.

There are numerous parameters controlling the molecular beam. Several, relating to the source and ablation laser (Nd-YAG), can be controlled by the computer. These are optimised by ‘scanning’ the parameter space - discussed in section 3.2. Additionally the source and ablation laser that must occasionally be realigned. The source is be aligned to the skimmer by monitoring gas pulses using a fast ion gauge (FIG) placed close to the PMT. In addition, the FIG should be used to check the carried gas pulse. profile, especially is the source seems noisy. If the FIG signal appears unstable the valve driver voltage should be reduced.

After a period of intensive use, the source becomes noisy and less intense. Upon inspection the target at this point is often pitted and tarnished, though very little ytterbium is actually consumed. Although it is straightforward to construct a new target, we find that resurfacing with a lathe, and thorough cleaning with *Scotch-brite* abrasive cloth provides good results. The target is cleaned with acetone in an ultrasonic bath before being installed in the chamber.

3.1.1 The time of flight (TOF) data

At present, for creating a molecular beam, we are solely interested in the probe PMT LIF signal. Each single TOF is a probe of the $|F = 1\rangle$ state population, passing through the probe beam (as well as scattered probe light, dark counts and electrical noise). An example is shown in figure 3.1. The raw units of acquired data are volts, though it is possible to converted to photon count rate⁵. This TOF is the signal observed from a single shot. Evident here is the distribution of the (YbF) molecular pulse, due primarily to thermal distribution of the molecules. From the TOF we can calculate the temperature of the molecular pulse. Though we would hope that all molecules undergo the same experimental conditions, this is an approximation, and on occasion we slice the pulse into velocity classes using the arrival timing information⁶, for example, to investigate spatial dependance of rf transitions. Also evident here is noise (primarily probe scatter) in the wings of the TOF. The time axis is simply the time from when the Q-switch fires to detection.

We can use this signal in two ways. Most commonly, we integrate over the TOF signal for a single shot. This gives us a single number proportional to the total number of YbF molecules

⁴see section 3.7.1 for a discussion of the shot noise limit.

⁵The conversion, measured by Condylis [14], is 0.148 V μ s/photon count.

⁶See sections 5.1 on field mapping and 4.2.3 on systematic effects.




Figure 3.1: Example time of flight (TOF). As the pulse of YbF molecules pass through the resonant probe laser LIF is used to detect molecules in the $|F = 1\rangle$ state. For most optimisation scans we integrate this signal to yield one number for each shot - the integrated TOF (in units of $V\mu s$.)

detected in the $|F = 1\rangle$ state. For a few tasks we use the timing information in the pulse to infer spatial information, so might split the TOF into bins and analyse these independently. This is covered in detail in the field mapping section 5.1.

3.2 Scanning parameter spaces

The next experiment we should consider is to explore a parameter space, an essential step in optimising the source, rf systems, and in electromagnetic field mapping diagnostics. To perform a scan, multiple shots are recorded, successively incrementing the relevant parameter. In the most simple case, the parameter space is 1-D, so we need just increment one parameter.

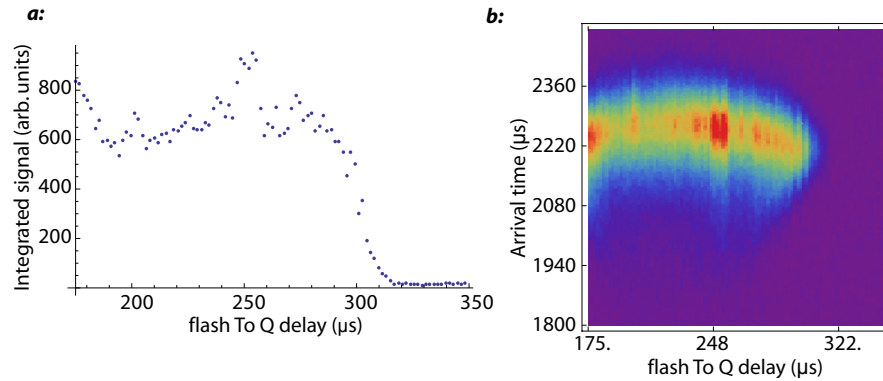


Figure 3.2: Scan of *flash-to-Q* timing (τ_{FtQ}), the time delay between pumping the Nd-YAG and firing the Q-switch. This parameter is used to control the ablation laser power. *a*: The integrated TOF curves most commonly used to optimise a parameter. *b*: The 2-D data, showing the timing/parameter information. This might be useful if we were aiming to create a slow beam, but is typically only used in mapping experiments.

Figure 3.2 shows a typical scan, in this case the source parameter *flash-to-Q* (τ_{FtQ}). This is the timing delay between the Nd-YAG pump flash, and the Q-switch firing. Tuning this parameter is our method of controlling the ablation laser power. A simple task might be to maximise the LIF signal. Here we fire a series of shots, with increasing τ_{FtQ} . We then plot the integrated TOF

against the parameter value, as shown in 3.2.a. Generally, the integrated TOF is the most useful tool, as we want to maximise the signal over all velocity classes. If we were interested in creating the slowest beam, for example, it would make sense to view the 2-D TOF data shown in 3.2.b, which shows the changing distribution of arrival times as a function of τ_{FtQ} .

This procedure can be used for all source and rf parameters. In practice, taking EDM data requires running the beam for long periods of time without optimisation. The biggest signal on a fresh target does not necessarily indicate the optimal configuration for taking EDM data. Optimal parameterisations vary as the ytterbium target ‘ages’. For τ_{FtQ} , tuning to a value slightly lower than the sharp decay of signal - approximately 275 μs in figure 3.2.a. tends to give good long term stability.

Other source timing parameters include the timing of the gas valve, relative to the Q-switch (*valveToQ*), and the length of time the valve is open for *valvePulseLength*. Optimisation of these parameters is carried out in a similar fashion. A full discussion of this source is presented in Condylis’ thesis [14]. Additionally, a table of typical parameterisations is included in appendix A.1.

3.3 State preparation

The first stage of the interferometer is to prepare the molecular pulse by depleting the $|1, \pm 1\rangle$ states. This is achieved by pumping the $|1, \pm 1\rangle$ states to the next electronic excited state ($A^2\Pi_{1/2}(\nu = 0, N = 0)$) using the Spectra 380 dye laser, resonant on the $Q(0)F=1$ transition (for molecular structure refer to figure 1.9). From this excited state, the molecules largely decay to states which no longer play any part in the experiment. However, the pumping is not complete - some molecules do decay back to the $X^2\Sigma^+(\nu = 0, N = 0)$ $|0\rangle$ and $|1, \pm 1\rangle$ ground states. These appear as DC background in our interferometer. The molecules then pass through the interaction region to the probe beam, where we drive an identical transition in order to measure the $|1, \pm 1\rangle$ state population using LIF. The ‘prepared’ molecules however are nominally in the $F=0$ state, so we expect no signal to be detected. To optimise the pump laser, the source is run in *align* mode (as above), which fires a series of shots, without changing any parameters. The operator can then adjust the pump laser alignment to minimise probe region LIF signal. Typically we observe a pumped signal of around 10% the unpumped amplitude.

3.4 A single resonance rf experiment

Having optimised the source and discussed scanning of parameters, we move on to a single rf transition experiment.

In order to test a single rf transition we just repeat the above experiment - pumping out the $F=1$

state - but insert an rf π -pulse between the pump and probe stages. The π -pulse requires applying an rf field, of precisely the correct length, amplitude and polarisation to drive the desired transition. The rf field is linearly polarised along the y -axis, which can be described as a superposition of right and left circularly polarised fields aligned to the z -axis⁷ (to which the molecule is quantised), such that we drive σ^+ and σ^- transitions, coherently populating the $|1, \pm 1\rangle$ sublevels.

The first parameter to consider is the rf pulse length. In Hudson’s thesis [16], long quasi-DC pulses were used. The molecules passed through a short region of homogeneous electric field. Given the relative stark shift of the $|0\rangle$ and $|1, \pm 1\rangle$ states, the application of carefully tuned rf would lead to a well defined π -pulse. This approach brings disadvantages. Firstly the linewidth is affected by inhomogeneous Stark broadening, which limits the efficacy of the π -pulse. Off resonance the rotation cannot drive a full π -pulse, and interferometer contrast is lost. Secondly driving rf transitions in spatially varying electric fields leaves us susceptible to fringe field systematic effects.

We have now moved to a short pulse regime. Still using a three region electric field design, the distance the molecules move during the rf pulse is much less than the length of the homogeneous field. The rf switches have a rise time of a few microseconds, which is small compared to the length of the rf pulse. There is no obvious systematic that arises from using arbitrary (yet repeatedly so) shaped rf pulses, but to calculate a closed form lineshape, we assume a rectangular amplitude envelope to simplify the mathematics. We usually run with 18 μ s pulses, though for certain task such as measuring Zeeman splitting we use longer pulses to reduce the transition linewidth.

The next parameter to set is the rf pulse timing. Assuming the molecular pulse velocity is constant, the pulse timing controls the point in space where the rf transition takes place⁸. For the EDM experiment care must be taken that this is away from the end of the plates where there are electric field gradients. Also, if using rf loops, it is important to drive transitions in the centre of the loops, where we maximise the number of molecules that receive a π -pulse. Furthermore, some regions are to be avoided near the ends of the apparatus where magnetic fields are not well controlled.

Under ordinary running conditions we apply an electric field of between 3.3 kVcm⁻¹ and 20 kVcm⁻¹. The relative Stark shifts of $|0, 0\rangle$ and $|1, 1\rangle$ levels are known from theoretical calculation, as well as previous experiments [43]: we can predict where the transition will be found (see appendix C). To locate the centre of the rf transition precisely we scan the rf frequency over the transition, as given a certain plate potential, we certainly know where the transition will lie to within the transition linewidth. Typically we have a good idea of the rf power required for a full π -pulse, though to avoid power broadening we might choose to underpower. To find the transition frequency we then fit to the theoretical lineshape, equation 3.1, for the rf transition (derived in

⁷Recall, the \hat{z} -axis is nominally defined by the electric field vector, whilst the \hat{y} axis is defined by the direction of propagation of the molecular beam. Refer to figure 2.1 for orientation.

⁸The distribution of molecular velocities, as well as uncertainties in the pulse origin should be taken into account if position information is important. For example see §5.1.4.

appendix B),

$$P_{F=1} = |C_q(t)|^2 = \frac{4b^2}{(\omega_0 - \omega)^2 + 4b^2} \sin^2\left\{\frac{1}{2}[(\omega_0 - \omega)^2 + 4b^2]^{\frac{1}{2}}t\right\}, \quad (3.1)$$

$P(F = 1)$ is the probability of observing the molecule in state $F = 1$ which was prepared in state $F = 0$. The frequencies ω_0 and ω are the transition frequency and applied rf frequency respectively, t is the duration of the interaction, and b is the Rabi frequency, a measure of the interaction strength.

A scan of rf frequency is shown in figure 3.3. Figure *a* shows the integrated TOF data, fitted to the rf transition lineshape described in equation 3.1.




Figure 3.3: Scan of rf field frequency. *a*: Integrated TOF scan. This scan, taken well inside the electric field, is a good example of the rf lineshape. The brown curve is a fit to the theoretical lineshape - equation 3.1. *b*: The raw TOF data shows little additional structure, other than the gaussian envelope of the molecular pulse in the time axis. However, fits of different velocity groups can be used to infer electric field values across the pulse 5.1.2.

The final task is to tune the rf power. Having set the rf frequency, and fixed the rf pulse length, we can now scan the rf power over a suitable range to observe more than the required π transition. We then set the power to deliver exactly π pulse, the first peak of the Rabi flopping curve. For the EDM experiment, where two pulses are required, each rf transition must be separately optimised. Figure 3.4.a shows a plot of the integrated TOF signal. This clearly shows as increasing rf power is applied, the transition probability reaches a maximum before oscillating with increasing frequency. We expect the rf transition probability to vary as $\sin[bt]$, where b represents the interaction strength, parameterised here by the rf field amplitude, and t is time. Expressed in terms of power, this expression becomes $P_{|1,\pm 1\rangle} = \sin[10^{b/10}t]$. Figure 3.4.a is a 3-dimensional plot of all recorded TOF data. In this data, taken using rf plates rather than loops, each velocity group received the same rf power, so the rf transition appears homogeneous across space (time axis).

When using rf loops, with optimal timings, molecules in the wings of the molecular pulse receive less rf power than those with mean velocity in the centre of the pulse. In figure 3.5, a scan of rf power taken using rf loops with 18 μ s rf pulses. As we scan the rf power, molecules at the centre of the pulse are transferred to the $|1,\pm 1\rangle$ states at lower powers than those in the wings. On this

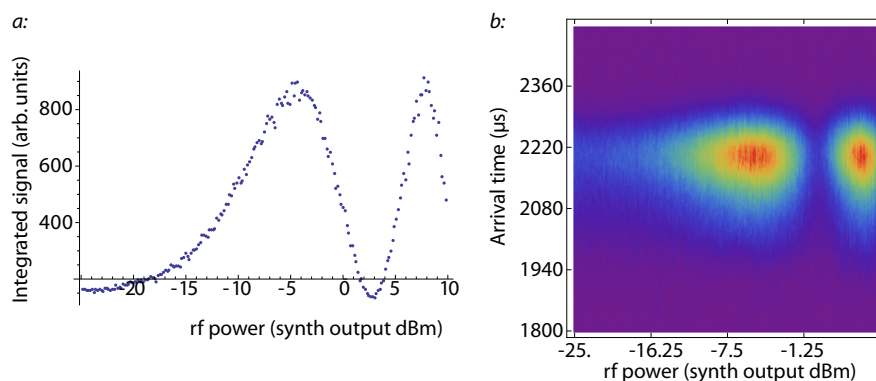


Figure 3.4: Optimisation of rf power using the rf transmission line structure. A parameter scan is recorded with increasing rf power. *a*: The integrated TOF signal, shows Rabi oscillations (This can be plotted on a log scale to yield a sinusoidally varying amplitude). *b*: The raw TOF shows the oscillations in power appear in phase across all arrival times. This is interpreted as all velocity classes receiving similar powers, as we expect for the homogenous fields generated using the rf transmission line.

occasion the integrated TOF loses some information, as molecules in different velocity classes trace out different Rabi oscillation curves. The result of this is that monitoring the integrated signal, the oscillations appear to wash out after the initial π -flip maximum.

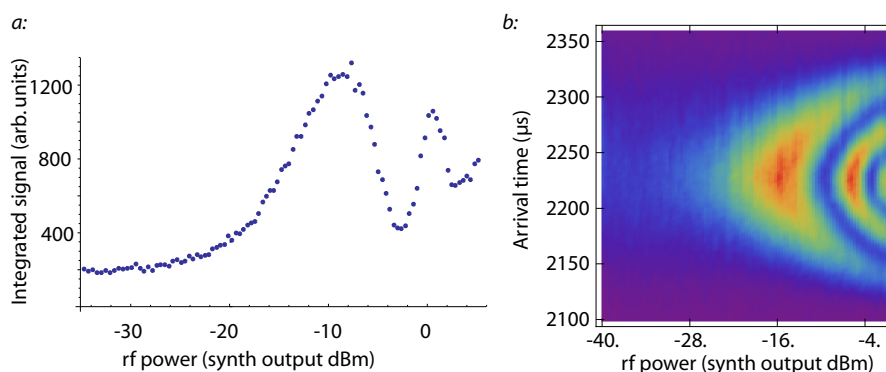


Figure 3.5: Scan of rf power using rf loops instead of rf plates. The field amplitude decays rapidly away from the plane of the loop, leading to an inhomogeneous rf transition across the pulse. This is most obvious in *b*: where the molecules in the wings of the molecular pulse (slow and fast) receive less rf power, and as such require a higher rf synthesiser power to achieve a π pulse. It is subtly apparent in *a*: also, where the Rabi oscillations appear to diminish with increased rf power, as a result of integrating over different velocity classes.

There are some subtleties in the rf systems, arising from the choice of using the rf transmission line, or rf loops. Using loops, it is only possible to drive transitions at, or very close to the loops. Using the transmission line, transitions can be driven at any point along the beam line, within the plates, and for a short distance outside the plates, though in an inhomogeneous field. Inside the plates (transmission line), the rf field is sufficiently homogeneous that the rf amplitude across the molecular pulse is largely uniform, so all molecules can receive a π -pulse.

3.5 A double resonance rf experiment

Our aim here is to implement the interferometric technique described at the end of chapter 1, albeit in a manner that is not optimised for measuring an EDM. Recall the output to the interferometer is sinusoidal, dependant on phases due to both Stark and Zeeman interactions:

$$R_{\text{PMT}} = R_0 \sin^2(\phi_{d_e} + \phi_B). \quad (3.2)$$

The EDM phase is known to be small, ($\phi_{d_e} \ll \phi_B$), so we generate interference fringes by scanning the magnetic field.

Before attempting the experiment the single rf transitions must be optimised. Using the approach described in the previous section, each is tuned to deliver a π -pulse, appropriately placed at either end of the interaction region.

A molecular pulse propagating towards the interaction region is first prepared by having the $|1\rangle$ state pumped out⁹. The molecules now enter the interaction region (applied E and B). The first rf pulse creates the $\frac{1}{\sqrt{2}}(|1, 1\rangle + |1, -1\rangle)$ superposition. The rf pulses are typically 1000 μs apart, over which time the superposition is free to evolve. The second rf pulse moves some of the amplitude back into the $|F = 0\rangle$ state, before the $|F = 1\rangle$ state is probed. The experiment is as described in chapter 1, figure 1.8. If we now integrate over the TOF, this shot results in one data point. A series of shots are acquired over a range of applied magnetic fields B_z trace out the interference fringes.

The resulting interference fringes have a pitch characteristic of the magnetic field applied, whilst their phase can be used to estimate the projection of the lab magnetic field vector onto z , as defined by the electric field. An example of such an interference scan is shown in figure 3.6. A series of shots were recorded, whilst a voltage controlled current source was used to scan $\pm 5 \text{ mA}$ ($\pm 80 \text{ nT}$).

This experiment is performed routinely to calibrate the magnetic field bias required to centre the interferometer fringes when taking EDM data.

⁹Outside of the electric and magnetic fields all three $F=1$ sublevels are degenerate, therefore all are on resonance with the pump beam.

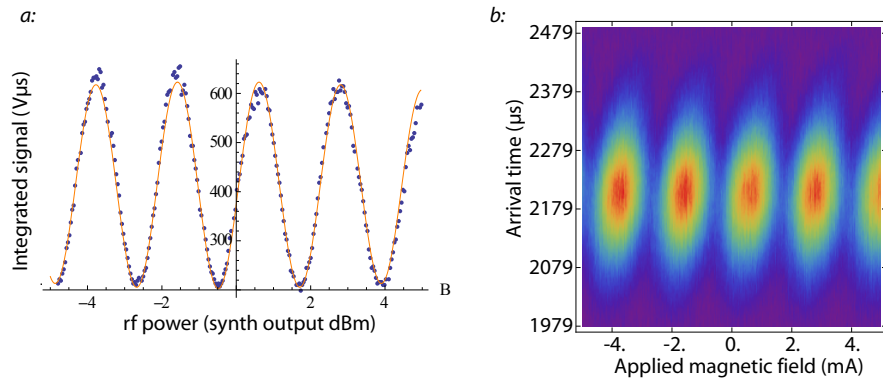


Figure 3.6: Interference scan of applied magnetic field B_z , using double pulse rf spectroscopy. The fit shown in figure *a*: is the expected theoretical lineshape $R_0 \sin^2[t\mu_B(B_{amb} + B_{appl})/\hbar + \phi_{de}]$. In this case the fringes have a phase offset of $159 \mu\text{A}$, equivalent to an ambient field of $B_z = 2.4 \text{ nT}$. Figure *b*: shows the full set of TOF data. Interestingly here there is a slight tilt to the fringes - a phase shift across velocity classes. This is symptomatic of field gradients near the end plates of the shield. It is also apparent from the plot that the amplitude of the oscillations appears to decay. This is due to the source decaying. For this particular scan, the first shots were taken to the left. It is common for the source intensity to decay significantly over the first few minutes, so single scans often display some degree of asymmetry. This is the motivation for the slow bit of each EDM data block, making the switching patterns symmetric about the centre (see section 3.6.3). Any systematic linear drift in signal is rejected in this manner. This is also the reason why despite some apparent redundancy within a single block, we consider a block to be the smallest unit of data to express a meaningful EDM.

3.6 EDM Experiment aspects

3.6.1 Measuring the EDM

Having constructed a working interferometer we are ready to make an EDM measurement: measuring a phase shift of the fringes under applied electric field. In a simple experiment, we could adjust the magnetic field B , so that the interferometer phase is halfway down a fringe, at $\phi_I = -\frac{\pi}{4}$. From this point we could repeatedly switch the direction of the E field. A measurable EDM would impart a phase on the fringes, observable as a change in the integrated LIF signal amplitude. With sufficient integration time, we would expect the observed LIF signal to be correlated with the direction of the electric field.

This approach is a form of phase sensitive detection. We can think of this as controlling the electric field polarity using a square wave. Call the control wave $W = + - + - + - + -$, which we use to switch the electric field polarity. We then construct an experiment to record a set of 8 shots, $S = s_i$. As successive shots are taken we increment through the control wave W , switching the electric field accordingly. To analyse the data we group the shots into similar states. In this case, with one switch, there are two *states*, so we have two sets of TOFs; $E^+ = \{s_1, s_3, s_5, s_7\}$, $E^- = \{s_2, s_4, s_6, s_8\}$. The EDM analysis is then just the mean difference in LIF observed in each of the two sets.

$$\phi_{d_e} \propto \frac{1}{8} \left(\sum E^+ - \sum E^- \right). \quad (3.3)$$

The problem with such a scheme is that we are sensitive to many other effects. We are implicitly assuming that nothing else in the beam or lab environment is changing. It is only possible with this approach to measure the phase of function $R_{\text{noise}} + R_0 \sin^2(\phi_B + \phi_{d_e})$, if the noise, scaling R_0 , and the phase ϕ_B are *constant*. This is clearly a poor assumption; we need a more sophisticated method. Additionally, the difference between the sums is a voltage. In order to convert this to a phase (and, in turn, an EDM) we would need to know the gradient of the fringes at $-\pi/4$.

Taking the other extreme, we could repetitively scan the entire interference lineshape (as in figure 3.6) under different electric field polarities. This would provide enough detail to extract information on phase shifts, interference fringe frequency, amplitude and DC offset, but we would record a lot of data at the peaks and troughs of the fringes - where we are insensitive to phase shifts.

So, next we consider stepping from one fringe to the next, recording 4 measurements $(\pm B, \pm E)$. Our analysis would now be to measure the phase of the fringes with two points. This approach is much more robust and can now distinguish a change in R_0 , the source intensity say, from a change in ϕ_{d_e} . Labelling the four states of the experiment $\{B^+ E^+, B^- E^+, B^- E^-, B^+ E^-\}$ we look for an

EDM Through the summation:

$$\phi_{d_e} \propto \frac{1}{4} \left(\left(\sum B^- E^+ - \sum B^- E^- \right) - \left(\sum B^+ E^+ - \sum B^+ E^- \right) \right) \quad (3.4)$$

Finally we need to measure the slope of the fringes, so we can convert a change in PMT count rate to an EDM. We add a small magnetic field step, ΔB . The size of this step is chosen to be small enough that the additional field leads to an approximately linear change in LIF signal, whilst large enough that we can measure the gradient accurately. This gives the experiment 8 states that are all combinations of $\pm E$, $\pm B$ and $\pm \Delta B$.




Figure 3.7: Measurement of the EDM phase, through a series of 8 interference experiments. Each circular point is an integrated TOF. Measurements are taken at 4 magnetic fields in each of two electric field states, in effect measuring the interference curve in each E polarity. Points $A - D$ describe one interference curve, $E - H$ the other.

For a minimal EDM measurement we record a TOF in each of the 8 states. The 8 integrated TOF signals¹⁰ now describe the two interference curves, in each state of E . The magnetic field steps used are $\Delta B = 100 \mu\text{A} \equiv \pm 1.6 \text{ nT}$, and $B = 560 \mu\text{A} \equiv \pm 8.96 \text{ nT}$. The magnitude of B is chosen to match $\phi_I = \pi/4$.

This stepping approach also allows for important diagnostic information to be extracted simultaneously. For example, if we take the sum of the 8 integrated TOF profiles, we have a direct measurement of the total signal integrated, which may be useful for monitoring source performance. Similarly compare the total signal integrated in each of the magnetic field states $\pm B$. Assuming the size of this step is correct, we can infer from this measurement how well we are centred on the interference fringes. We refer to each one of these calculations as *analysis channels*. These ideas

¹⁰It is certainly possible to bin the data by arrival time to search for a varying EDM or channel analysis across the pulse, and can be useful for identifying systematics. This is discussed in section 4.2.3.

are developed more fully in the following section.

This scheme is not perfect however. Despite describing the interference curve with 4 points, there can still exist shifts in the lineshape which we interpret as a phase shift due to the EDM which are not. This systematic effect is analysed in section 4.2.2.

3.6.2 Analysis of the EDM measurement

Analysis of the interferometer lineshape is simplified by approximation to a first order expansion about $\phi_B = \pm\pi/4$.

The interferometer lineshape can thus be approximated about $\frac{\pi}{4}$ to a piecewise linear function:

$$R = R_0 \sin^2(\phi_I)$$

$$\simeq R_0 \begin{cases} \frac{1}{2} - \frac{\pi}{4} - \phi_I & \phi_I < 0 \\ \frac{1}{2} - \frac{\pi}{4} + \phi_I & \phi_I \geq 0 \end{cases},$$

where $\phi_I = \phi_{de} + \phi_B$. A comparison of the approximate and exact theoretical lineshapes is shown in figure 3.8.




Figure 3.8: Comparison between the theoretical lineshape and the simplified, piecewise linear lineshape. The blue markers are the magnetic field points $\pm B \pm \Delta B$.

In order for the large magnetic field step B to fall very close to $\phi_B = \pm\pi/4$, the ambient DC magnetic field must be trimmed as even with shielding we see some residual lab field. Given a magnetic shielding factor of 1000, a lab field of 30 μT (approximately the earth's magnetic field) is reduced to 30 nT at the interaction region. This would lead to an interferometer phase of approximately 0.2 rad. The concern is not that this will induce an EDM; a static field generally¹¹ cannot do so, as it must be modulated in phase with $B.E$. Rather, the DC field will shift the fringes such that the B step is not perfectly centred upon a fringe. In this case the assumption that the lineshape is linear close to $\pi/4$ fails. Assuming the lab field vector is homogeneous, and the applied magnetic field too is homogeneous, we can null the DC component of the lab field along B_z by trimming the current source used for magnetic field steps. To do this we scan the interference fringes, and measure the displacement of the centre fringe in units of current, then offset

¹¹There are non-obvious effects that can make a static B_y field look like an EDM - see section 4.2.1.

all magnetic field steps appropriately. Here we have cause to use the analysis channels previously introduced in section 3.6.1. Whilst taking EDM data we can use the ‘ B -channel’ (change in signal correlated with the B step from one fringe to the next) to watch for drifting ‘static’ magnetic fields. Assuming perfect trimming, this should average to zero. A small value of the B -channel analysis channel means we are operating close to $\pi/4$ on the interference curve. Ambient magnetic fields of frequency much greater than the pulse rate of the experiment are highly attenuated by the shielding, and in any case will time average and not contribute to the EDM signal (though the may prove a source of noise).

Another useful analysis channel is the E -channel - the change in signal correlated to a change in the electric field polarity. For this we take the mean of all points where the electric field is in one orientation, and subtract the mean of points in the other orientation. It may appear as if the change in signal correlated with E should be the EDM signal. This is not the case due to the change in gradient of the lineshape around $\phi = 0$. The EDM analysis must sign with the B channel, and is in fact the product $E.B$. The E analysis reflects the situation where the lineshape has better contrast in one E -state, an indication that an electric field reversal asymmetry leads to more efficient rf transitions in one polarity.

Table 3.1 summarises the main analysis combinations used.

Channel Combination	Comment
Total Signal $\frac{1}{8}(A + B + C + D + E + F + G + H)$	The mean signal observed
E shift $\frac{1}{8}(A + B + C + D - E - F - G - H)$	Intensity shift, correlated to E reversal
B Shift $\frac{1}{8}(A - B - C + D + E - F - G + H)$	Intensity shift correlated with B . How well the magnetic field is nulled.
CAL , (ΔB shift) $\frac{1}{8}(A - B - C + D + E - F - G + H)$	Change in signal due to a small change in ΔB . The gradient of the fringe
EDM $\frac{1}{8}(-A - B + C + D + E + F - G - H)$	The EDM channel. Change in signal with change in $E.B$
$E.Cal$ $\frac{1}{8}(A - B - C + D - E + F - G - H)$	The change in slope with change in E
$B.Cal$ $\frac{1}{8}(-A + B - C + D - E + F - G + H)$	The change in slope with change in B

Table 3.1: Analysis of stepped interferometer scheme data.

The *Total signal*, is simply the mean signal observed. It primarily it is used to monitor source quality¹². The EDM signal is correlated with E and B ($E.B$). The B -shift is the change in signal correlated with the large B step, $\phi_B = -\pi/4 \longrightarrow \phi_B = \pi/4$. A non-zero value in this channel indicates that the fringes are not centred accurately about $\phi = 0$. The final important channel is CAL , the change in signal associated with the small ΔB step. This step is much smaller than

¹²The EDM software, *Scanmaster*, also offers a real time full TOF view which is also useful for this task.

the fringe period. It measures the change in signal with a change in ΔB - the gradient of the fringes. Using this gradient, and knowing the size of the small ΔB step, it is possible to convert other interferometer signals from count rates to physical units. The correlation of CAL with the electric field state, $E.CAL$, is used to cancel a systematic artefact (§4.2.2) of our linear lineshape assumption. The $B.CAL$ and $E.B.CAL$ channels are more complicated correlations, not commonly used in our analysis.

3.6.3 Development of the EDM measurement scheme

In the previous section a scheme was introduced which modulates three steps, to give 8 machine states. In this section I develop this scheme into a more general method, and introduce clearer notation.

In the simplest sense we previously considered combinations of $\pm E$, $\pm B$ and $\pm \Delta B$ to iterate systematically through the 8 states. To do this each quantity is switched between two states, which can be logically considered on/off, using a switching waveform[45].

To generate the switching waveforms, define a square wave of period T , termed 1. A square wave of period $2T$ is termed 10, and so on. We can now control the state of E , B and ΔB using these three waveforms.

$$\begin{aligned} W_E &= 001 = - - - - + + + + \\ W_B &= 010 = - - + + - - + + \\ W_{\Delta B} &= 100 = - + - + - + - + \end{aligned} \tag{3.5}$$

As before, to analyse the data for a signal correlated with B , for example, we take the mean of points where the W_{B_i} is '+', and subtract the mean of points where W_{B_i} is '-'. A key advantage of this approach is in describing analysis channels correlated to two patterns, like the EDM channel. This amounts to taking the XOR each pair of binary digits. With this compound waveform, the EDM analysis is the sum of shots where $W_{E.B}$ is positive, minus the shots where it is negative.

$$W_{E.B} = W_{B_i} \text{ XOR } W_{E_i} = 110 = - - + + + + - - . \tag{3.6}$$

In practice this approach is modified a little. First of all, we use a basis set which contains longer waveforms, so that E is switched very slowly¹³, compared to the B and ΔB . The switching waveforms have 12 bits, making each waveform $4096 T$ long. The electric field is switched at $W_E = 110000000000$. We use the 2nd bit to make the E reversal pattern symmetric with respect to the block, which rejects systematic effects caused by slow drifting magnetic fields, appearing similar in form to the E - waveform shown in 3.9. Faster switches are similarly switched using products of

¹³It would be preferable to switch everything fast to move away from $1/f$ noise, but technical reasons concerning the fast switching of electric fields (see §3.7.2) prevent this.




Figure 3.9: Example switching and analysis waveforms.

square waves. The large magnetic field step is switched according to $W_B = 000000011011$, whilst the small ΔB field step is switched at $W_{\Delta B} = 111100110111$. The EDM analysis therefore is $W_{\text{EDM}} = 110000011011$.

So far we have concentrated upon the minimal set of channels required to measure an EDM. However the 12 bit waveforms allow us to introduce many extra channels for probing systematics. For the EDM data presented later in this thesis we modulate only one additional parameter, a phase shift between the first and second rf pulses. Whenever the waveform is ‘1’ we shift the phase of the second pulse by 180° , when 0 we introduce no phase shift. The aim of this is to identify any coherence between the $|0\rangle$ and $|1\rangle$ states which might result from imperfect rf transitions, and could lead to a systematic EDM. In the case that the rf transition is partial - say a $\pi/2$ transition - then the $|F=0\rangle$ and $F=1\rangle$ states will be coherent, apparent in the interferometer as a Stark shift dependent phase. Using the full rf transition matrix (appendix B) it is straightforward to calculate that shifting the phase of the second π - pulse on alternate shots, the net $|F=0\rangle - |F=1\rangle$ phase is rejected from the interferometer lineshape. In contrast, analysis of the π -flip channel can expose *only* this contribution. If the rf transitions are perfect π pulses, this phase is necessarily zero, and hence the π -flip has no effect on the interferometer.

Uncertainties

Error bars figure prominently in the discussion of an EDM. Using our waveform switching scheme, I have described how to extract central values by taking linear sums of integrated TOF’s. However, I have not yet addressed how we can estimate error bars.

In fully executing one 12-bit waveform, we switch 4 parameters ($B, \Delta B, E, \pi$ - flip) over 2^{12} shots, which we call a *block* of data. However, the number of different combinations of the 4 switches is $2^4 = 16$ states. Therefore each block contains 256 shots in an identical state. These are the sub-blocks alluded to in chapter 3. For each state, this makes a set of points s_i , with variances σ_i^2 . In general, our switching distributes these states throughout the block. The simplest analysis would be to firstly take means and variances of each set. The central value for each channel could

then be calculated by taking linear sums of the correctly signed means, whilst the error bar for each channel would be the sum of the variances $\sigma^2 = \sigma_1^2 + \sigma_2^2 \dots = \sum_i \sigma_i^2$.

This in fact is slightly too restrictive. The EDM analysis partitions the block into 16 subblocks corresponding to the 16 states. The points within each subblock are uncorrelated, independent, taken under identical experimental conditions. However, the subblocks themselves are somewhat correlated. For example, shots taken in states $+\Delta B$ and $-\Delta B$ are related by the gradient of the fringes. The variance of the sum of uncorrelated variables is simply the sum of the variances. The variance of a sum of correlated variances is the sum of the covariances:

$$\text{Var} \left(\sum_i X_i \right) = \sum_i \sum_j \text{Cov}(X_i, X_j), \quad (3.7)$$

where the covariance matrix $\text{Cov}(X_i, X_j) = E((X_i - \bar{X}_i)(X_j - \bar{X}_j))$ is the covariance of subsets X_i, X_j .

By considering the correlations between subblocks, we can calculate an appropriate variance for each analysis channel. The compilation of blocks into a single EDM measurement is discussed in chapter 6.

3.7 Noise

There are many sources of noise in our experiment; shot noise, magnetic field noise, source noise, probe laser noise and electrical interference. In the following sections I discuss the main sources of noise, and how they might effect an EDM measurements.

3.7.1 Detector noise

The most obvious source of noise in the experiment is the PMT. To give an indication of the experimental limit of our experiment we can perform an elementary statistical analysis, based upon a crude EDM experiment. This was originally carried out by Hudson ([16], p. 66), but is repeated here for clarity and to reflect experimental advances.

In figure 3.10 I show the interferometer lineshape, along with the linear lineshape approximation close to $\phi = \pi/4$. A non-zero EDM interaction would introduce some phase, which can be measured as a change in the integrated PMT current.

If we spend equal amounts of time with the field in each orientation, our net signal is:

$$S = \frac{T}{2}(R_a + R_b). \quad (3.8)$$

Around the point $\phi = \pi/4$, the count rate is $R_0(1/2 + \phi_{de})$, where $\phi_{de} = \eta d_e E_{\text{eff}} T / \hbar$. Counting for a period of time $T = t/2$ gives $N_{\pm} = R_0 \frac{t}{2}(1/2 + \phi_{de})$, with statistical noise $\sigma_{N_{\pm}} = \sqrt{\frac{1}{2} R_0 t}$, and the difference in count rate is $(N_+ - N_-) = \phi_{de} R_0 t = R_0 \eta d_e E_{\text{eff}} t / \hbar$. In general there is also

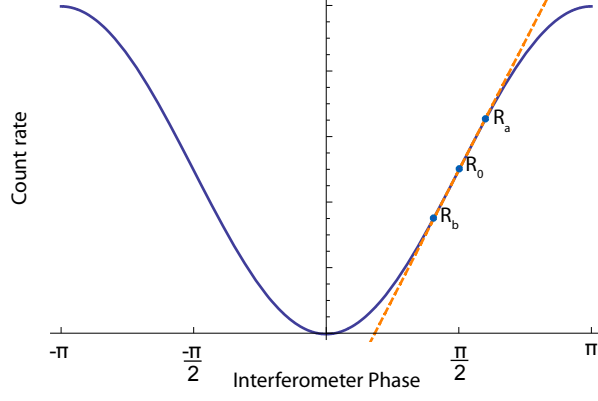


Figure 3.10: Illustration of simplest EDM measurement we could make. The straight line represents the assumption that the lineshape is linear close to $\phi = \pi/4$.

a background rate, R_B giving a total shot noise of $\sqrt{R_B t + \frac{1}{2} R_0 t}$. The statistical signal to noise ratio is then

$$S : N = \frac{\phi_{de} R_0 t}{\sqrt{\frac{1}{2} R_0 t + R_B t}}. \quad (3.9)$$

In the case that $R_B t \ll \frac{1}{2} R_0 t$,

$$S : N \simeq \phi_{de} \sqrt{2 R_0 t}. \quad (3.10)$$

Finally, setting the signal to noise as 1,

$$d_e = \frac{\hbar}{\eta E_{\text{eff}} T} \cdot \frac{1}{\sqrt{2 R_0 t}} \quad (3.11)$$

Taking $\eta = 0.7$, $E_{\text{eff}} = 26 \text{ GVcm}^{-1}$ and $T = 1 \text{ ms}$ we find,

$$d_e = \frac{3.6 \times 10^{-23}}{\sqrt{2 R_0 t}} \text{ e.cm} \quad (3.12)$$

Using the integrated TOF signal, we replace $R_0 T$ with the counts per block. The conversion factor is $1.48 \text{ V}\mu\text{s}/\text{photon}$. A typical integrated TOF is $650 \text{ V}\mu\text{s}$, thus we detect $4400 \text{ photons/shot}$. Given that a block of EDM data is 4096 shots , the correct substitution is $R_0 T \rightarrow 4400 \times 4096$. This yields a statistical limit of $8.4 \times 10^{-27} \text{ e.cm}/\sqrt{\text{block}}$.

This analysis of the statistical noise is a useful insight, as it ensures that our acquisition will not be limited to detector efficiencies. It also provides us a sensitivity to aim for when reducing other noise sources.

3.7.2 Source noise

Some of the total noise comes from non-statistical variations in the pulses of molecules produced by the source. As a brief investigation, 40000 TOF pulses were analysed. The data was gathered using the normalisation PMT which probes the initial $F = 1$ population. The TOF curves were integrated, then the set was Fourier transformed to a power spectrum (figure 3.11).




Figure 3.11: *a*: Time series of 40000 integrated, unpumped TOF signals. *b*: Power density spectrum of the unpumped TOF time series. The dark blue curve is a fit to $\alpha + \frac{1}{(f+\delta)}$. The offset delta is introduced to avoid fitting the singularity at 0 Hz.

From figure 3.11.*a*, it is clear that the source is qualitatively noisy, displaying occasional drops in signal. Using the power spectrum it seems reasonable to characterise the noise as having a $1/f$ component and a white noise component. We can see that the noise power is approximately $0.9 \text{ V}\cdot\mu\text{s}^2\text{Hz}^{-1}$ above 6 Hz where the power density spectrum is almost independent of frequency (white noise). At low frequency we see considerable $1/f$ noise. The key observation here is that we should strive to take EDM data away from the $1/f$ dominated end of the spectrum.

3.7.3 Magnetic field noise

Magnetic field noise can contribute to noise in the interferometer signal. There are two ways to group magnetic field noise, random and systematic. Random noise is such that, by the central limit theorem the mean phase imparted to the interferometer over many shots is zero. However, the uncertainties in magnetic field noise introduce uncertainties to the measured EDM.

We should also be mindful of magnetic field ‘noise’ which, though perhaps apparently random in nature, has a component which switches with the W_E waveform. This mimics the EDM exactly¹⁴.

Using a flux gate magnetometer positioned outside the inner shield we probe the magnetic field whilst taking EDM data. Several magnetic field measurements are taken for each shot, at times when the molecules are in the interferometer region. These points are then integrated to provide an integrated magnetic field for every shot. The resulting 40000 point time series was then expressed

¹⁴The EDM in the PMT data is characterised by a correlation to the $E.B$ channel, whilst in the magnetometer it is correlated to just the E channel. This is because the interference lineshape gradient changes sign at $\phi_I = 0$. Considering the fringes it is clear that to simulate an EDM with a magnetic field we just need to shift them some small angle whenever E reverses. Looking at the PMT analysis however, we need to do this and account for the points either side of $\phi_I = 0$ moving in opposite directions, hence the PMT edm analysis is $E.B$

as a power density spectrum, as shown in 3.12.




Figure 3.12: *a*: time series of 40000 integrated, magnetometer ‘shots’ *b*: Power density spectrum of the magnetometer time series.

This figure shows that the magnetic field noise is more complex in structure than the source noise. Similarly there is a $1/f$ type dependence at low frequency, there are higher frequency components at approximately 5.5 Hz, 10.75 Hz and 11.75 Hz. Whereas for the source noise we can draw some conclusions through inspection of the power density spectrum, it is less appropriate in this case, as it is not obvious how this noise is rejected by the EDM switching pattern.

To quantify how magnetic field noise limits our sensitivity we apply our switching pattern analysis (substituting the integrated magnetic field noise measurements in place of the integrated TOFs). This tells us how the integrated magnetometer signal is correlated to a specific switching pattern - and gives us an error bar on this quantity. We would like in this case to find how the magnetic field changes with the E channel¹⁵.

The raw magnetometry signal, analysed against the E switching channel has a 1σ uncertainty of $36 \mu\text{V}/\sqrt{\text{block}}$. Given a magnetometer calibration of $10 \mu\text{TV}^{-1}$, and a shielding factor of 1000, this can be expressed as an uncertainty upon the mean magnetic field uncertainty throughout the interferometer region the machine, corresponding to $3.6 \text{ pT}/\sqrt{\text{block}}$. This is magnetic field noise, inside the machine, which our phase sensitive detection does not reject. Clearly this will affect the phase accrued by our EDM interferometer. Expressed as an EDM, the magnetic noise is equivalent to $9.3 \times 10^{-27} \text{ e.cm}/\sqrt{\text{block}}$, or $4.2 \times 10^{-28} \text{ e.cm}/\sqrt{\text{day}}$. This suggests if we were limited solely by magnetic field noise we ought to be able to make a competitive EDM measurement in several hundred blocks.

Non gaussian statistics and the bootstrap

It was implicitly assumed in calculating the magnetic field noise above that the sampled points were normally distributed. The Fourier transform of such a signal would be flat. However, we know this is not the case, as from figure 3.12 we can see a rise in noise at low frequencies, as well

¹⁵Recall, the EDM channel analysis is $E.B$. As there is no change of sign with B , as was the case with the PMT data, we can just use E .

as the several peaks in the spectrum. This raises questions as to the validity of assuming gaussian statistics, and how best to report the EDM 68% confidence limit.

A useful approach is to use a statistical technique called the *bootstrap*, which I now describe. We start with a set of blocks containing data either acquired with the magnetometer (if investigating field noise), or the PMT (if measuring an EDM). Taking the case where we are interested in magnetometer data, we would analyse each block according to the E channel to extract the magnetometer EDM, as above. The result would be a set S of analysed block values b_i . In previous sections, we would take the list of b_i values, and generate a confidence interval using a weighted mean, assuming gaussian statistics.

The bootstrap [46] makes the assumption that the analysed blocks b_i are representative of some unknown parent distribution, and in fact offer the best description of this distribution. Using the bootstrap, a number of replicates of S are made. There are a number of approaches to this, the simplest of which is to draw randomly with replacement. An EDM is now calculated for each replicate set by taking a weighted mean of each block's EDM. The cumulative density function of the mean central value for each replicate set is now generated. From this we can calculate a 68% confidence interval which better reflects the 'width' of non-gaussian distributions.

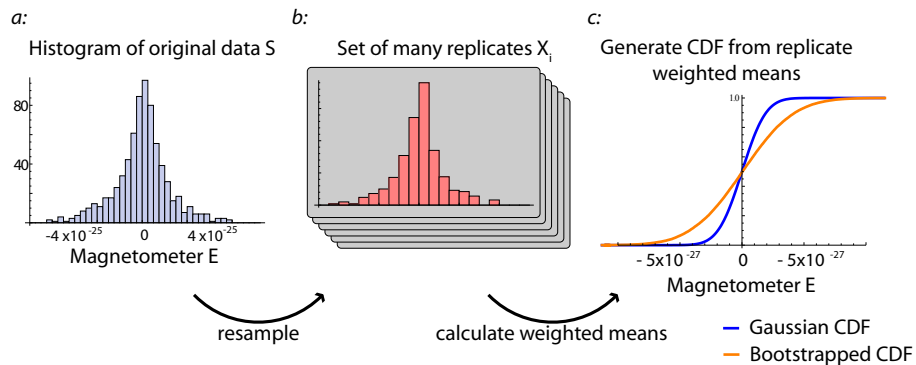


Figure 3.13: Illustrated bootstrap of the magnetometer E channel analysis. *a*: The histogram of single block results taken over 824 blocks. *b*: Set of 30000 replicate data sets. *c*: The central value of each replicate data set is used to generate a CDF (orange), which is compared to a normal CDF, whose variance and mean are calculated from the original data set.

A simple bootstrap is illustrated in figure 3.13. The input data is the magnetometer E channel data. As can be seen, the data is not normally distributed, with significant weight extending into the wings. From figure 3.13c it is clear that the bootstrap suggests a greater standard deviation than the simple variance of the original distribution. This reflects the extra weight in the wings evident in the original distribution 3.13.a. So, using the bootstrapped cumulative distribution function (CDF), we can calculate more accurately the effect of the magnetic field noise. The result, expressed as a 68% confidence interval, implies an uncertainty of $4.2 \times 10^{-26} \text{ e.cm}/\sqrt{\text{block}}$, or $1.9 \times 10^{-27} \text{ e.cm}/\sqrt{\text{day}}$.

The motivation for this analysis is that the distribution of both PMT and magnetometer EDM

analysis should not be assumed to be normally distributed. In section 6.1, in the EDM data analysis, we use the bootstrap to generate a 68% confidence interval that is clearly wider than the 1σ error bar calculated using the covariance method.

The magnetic field veto

The distribution of magnetic field noise has been shown not to be normally distributed. The additional weight in the wings of the distribution means the occupancy of ‘rare’ events is much higher than gaussian statistics would predict. The variance is perhaps not the best measure of a given distributions width.

Using gaussian statistics we know that taking more data will always improve our precision. If we record data where we assume the standard deviation to be σ , then we expect the standard error of the mean to scale as σ/\sqrt{N} , where N is the number of points samples.

In contrast, if the bootstrap is used to calculate confidence intervals, including the noisy wings of the data reduces overall precision. In the following analysis we introduce a veto to exclude the noisiest data. We then calculate the 68% confidence interval of the remaining data using the bootstrap.

To set a threshold, we compare the EDM error bar (PMT $B.E$), to the magnetometer E -channel. We then exclude blocks if the ratio of these two numbers is above some level.




Figure 3.14: EDM precision as a function of veto level. Successively increasing the veto level include more blocks into the analysis. The gaussian approximation demonstrates that additional data increases sensitivity. The bootstrap CI however shows that there is an optimal level at which to exclude noisy blocks.

The results, figure 3.14, show a veto level (~ 5), at which including more noisy blocks reduces the overall sensitivity.

3.7.4 Switching pattern noise rejection

Having discussed the switching patterns and noise spectra of our experiment it is useful to calculate the power density spectrum of our switching patterns in order to see how the two overlap. In the top row of figure 3.15 I show the power spectra of the four switching patterns, whilst four representative

analysis patterns are shown in the bottom row. These are to be compared with the source noise and magnetic noise spectra, figures 3.11 and 3.12.

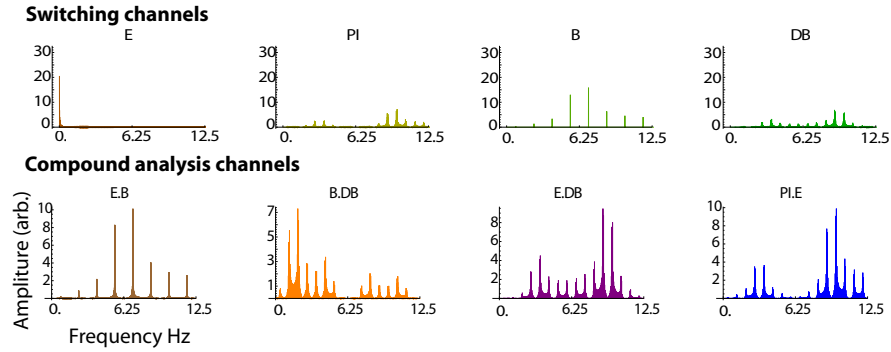


Figure 3.15: Power density spectra of switching patterns.

The less overlap with the noise spectra a pattern has, the better it will reject that noise¹⁶. We see that all patterns, with the exception of *E*, consist of components spread throughout the spectrum. The electric field cannot switch fast, and is forced to have frequency well below 1 Hz.

From this it is apparent that most channels are well configured reject both magnetic field noise and source noise, with the exception of the *E* channel.

3.7.5 Phase locking and mains noise

Magnetic field mains noise is a source of interference to our experiment. Though this averages to zero over many cycles, it can be removed altogether.




Figure 3.16: Moving-window, power density spectrum of magnetic field noise measured by the flux gate magnetometer sampling at 500 Hz. We recorded several hours of magnetic field fluctuations. We see the dominant noise occurs at 50Hz, with components spread throughout the spectrum.

¹⁶This fact can be used to select optimal switching patterns for the most important channels.

To investigate magnetic field noise in the lab, the fluxgate magnetometer was placed outside the outer shield, and recorded several hours of data at 500 Hz. A moving window power density spectrum was then calculated from this time series (figure 3.16). From this is clear that the magnetic field noise spectrum is dominated by several large peaks at 50, 150 and 250 Hz.

In order to hold the magnetic field noise constant, we phase lock our experiment to the mains, so that running at 25 Hz, every shot takes place at the same point of the mains cycle, as the mains related field is stroboscopically frozen.

Chapter 4

Systematic effects

Clearly the primary aim of our experiment is to improve the accuracy of the electron EDM measurement. Partly this includes improving the sensitivity of the experiment, but as we do so we also have to ensure that our final EDM analysis is due solely to the EDM interaction. Any quantity which can mimic the EDM signal is termed a *systematic effect*. In the following sections I discuss several such recently considered effects, and our efforts to suppress them.

4.1 Review of the EDM measurement concept

The interferometer measures the splitting due to the Zeeman and EDM interactions. The EDM is observed through a change in the measured interferometer phase correlated to the electric field state. However, under our interferometer scheme, the correct analysis is that the EDM is proportional to a change in signal correlated with a change in the state of $E.B$. Any candidate for a systematic therefore must be correlated to the same modulation.

The origins of systematic effects lie in the real world implementation of the experiment. A ‘real’ EDM appears as change in the population of the $|1, \pm 1\rangle$ states as a result of the splitting caused by the EDM interaction. Our simple model, as described by the interferometer lineshape (figure 1.10) makes various assumptions which in practice are manifestly untrue. Crucially we have thus far assumed:

- perfect rf transitions
- homogeneous fields
- no noise
- a ‘simple’ interferometer lineshape with no DC offset or changes in contrast

The analysis of systematic effects quantifies how relaxing each of these assumptions impacts our measurement, and ultimately whether we can be certain the EDM measurement is truly due to

the EDM interaction¹. As the required level of sensitivity is increased, we must be ever mindful of the validity of assumptions made in our analysis: rf transitions are not perfect, crosstalk between signals cannot be perfectly isolated, etc. Here I describe several approaches used to probe the various sources of systematic errors.

4.1.1 Probing for possible systematic effects

Null tests In general, these entail running the experiment without the interferometer. For example, a simple test is to acquire data without a molecular beam, lasers or rf, leaving solely the PMT output and the switching electric and magnetic fields. In this case, we analyse the PMT signal (noise and some ambient scatter) as normal. The measurement of a non zero EDM phase would imply a systematic such as pickup of the E and B control signals². A second obvious null test is to connect a battery to the computer input that is usually used to sample the probe PMT - providing a DC signal with a small amount of noise. Analysis of this data probes effects that might arise from software issues (systematically failing to record each shot for example), aliasing issues upon analogue to digital conversion, electrical pickup or analysis errors. Null tests fortunately can be carried out quickly. Ordinarily, EDM data acquires for 200 μ s, per shot, at 25 Hz. Null test data can be acquired a factor of 200 times faster using very wide gates. Null runs taken using a lamp to scatter light into the PMT recorded a result of the equivalent to an EDM of $6 \pm 2 \times 10^{-32}$ e.cm in March 2006. These experiments will need to be repeated in future EDM experiments, as the apparatus has changed significantly since this date.

Polar molecules insensitive to the EDM interaction Another form of test would be to repeat the experiment using a second polar molecule, insensitive to the EDM interaction, but with identical magnetic structure. Clearly, in this case the EDM analysis should present a result consistent with zero³. If it does not we can be certain that systematic effects are present. Once the experiment is running efficiently, and all other systematics are understood we plan to acquire a dataset using CaF, where the effective electric field is so small that the experiment could have no hope of measuring a real EDM.

EDM data taken at different electric fields Another check is to repeat the experiment at different voltages. In the presence of a real EDM the interferometer phase must scale non-linearly, due to the form of the molecular polarisation, η , (figure 1.5). Ordinarily the EDM is calculated

¹A trivial example of such effects was demonstrated after moving from three region fields to single region fields (discussed in §4.2.1). In the single plate scheme the molecules are on resonance with the rf throughout the interaction region. The old rf switches offered -30 dB isolation. Though it was sufficient before, the -30 dB leakage into the interferometer after upgrading the electric fields gave rise to EDMs of order 10^{-24} e.cm. The solution was to use two switches in series, with a total of -110 dB isolation.

²Or perhaps an unlikely modulation of the room lights or scatter with $E.B$.

³Ordinarily we scale the EDM analysis channel by the effective electric field for the molecule. For an insensitive molecule this is 0, and therefore if we scale EDM analysis, our result is accordingly zero. Therefore we would look at the raw analysis channels $E.B$.

using

$$d_e = \frac{E.B}{DB} \frac{100}{560} \frac{\pi}{4} \frac{\hbar}{t\eta E_{\text{eff}}}, \quad (4.1)$$

where $E.B$ and B are analysis channels. Rearranging this, we can calculate the interferometer phase as

$$\phi_{d_e} = \frac{E.B}{DB} \frac{100}{560} \frac{\pi}{4} \propto d_e \eta. \quad (4.2)$$

As can be seen, in the case that the EDM is truly non-zero, a robust EDM measurement must consist of a set of measurements taken over a range of electric fields. These must then be shown to scale with η .

Detailed modelling of the experiment Systematic effects arise as the idealised description fails to describe subtle effects in the real world implementation. The simple model assuming perfect rf transitions and homogeneous fields cannot be used reliably with real world fields and machining tolerances. Therefore, the final stage of understanding systematic effects is to develop a more detailed model of the experiment. We are in the process of developing a more general model for the experiment, which will calculate rf transitions over arbitrary electromagnetic fields. Combining this with detailed FEMLAB models⁴ (derived from CAD schematics) it should be possible to search thoroughly for less obvious systematic candidates. A useful tool to improve such modelling is to use the actual observed electromagnetic fields as model parameters. In chapter 5 new techniques which allow us to map the electric, rf and magnetic fields throughout the interferometer region are discussed. Given these we could construct accurate models that reflect the true nature of the experimental environment.

4.2 Specific systematic effects

4.2.1 Systematic effects arising from the use of a multiple electric field regions

In this section I will discuss in detail a systematic effect that was detected in recent work, and efforts made to reject it.

Until August 2006, the electric field used for the EDM experiment comprised of three separate regions. At each end of the interferometer, there were low field ‘guard’ (g) regions, separated by the main ‘centre’ (c) region. The rf transitions were driven in the guard regions, whilst the centre portion provided the electric field for the bulk of the interferometer⁵. Molecules were placed in

⁴FEMLAB is software that can be used to numerically calculate electromagnetic fields.

⁵The reason for choosing this scheme is that at lower voltages, plate spacing variations, due to machining tolerances, gives rise to less broadening than at higher voltages. Conversely, we obviously wish to run the interferometer

the quantum superposition within the guard region, then pass through to the centre region. We realised however that under certain conditions a systematic may arise from the curvature of the electric field between the g and c regions.

Within the interaction region the molecules are aligned to the electric field, which in turn defines the z axis of the experiment. Whilst in the interferometer, the molecules are aligned such that the interferometer phase depends only on B_z and the EDM interaction. This gives the total interferometer phase:

$$\phi_I = 2(\mu_B B_z - d_e E_{\text{eff}} \eta) \frac{t}{\hbar}. \quad (4.3)$$

Under these assumptions, when we reverse the electric field, any change in interferometer phase is due solely to the EDM interaction. However, if the magnetic field B depends upon the state of E , clearly this condition will fail.

Such a situation can arise under certain plate voltage asymmetries, whereby the 0 V planes in the c and g regions are not the same. As molecules move from the g to c regions, the 0 V plane bends. To illustrate this the electric potential was calculated numerically for a crude model of the plates. This is shown in figure 4.1, where for clarity the asymmetry is greatly exaggerated by a factor of 1000 more than likely asymmetries. The molecules will, as they pass through the interaction region, rotate a small angle $\theta(y)$ to align to the electric field. As they do so the interferometer no longer solely sensitive to $B_z(y)$, but rather $B(y)' = B_z(y) \cos \theta(y) + B_y(y) \sin \theta(y)$. If the 0 V plane configuration varies for each state of E , then the projection of the magnetic field onto the alignment axis of the molecule will change, leading to a systematic EDM.




Figure 4.1: Numerical calculation of electric potential between the guard and centre field plates, nominally at ± 2 kV. In this graphic, a huge asymmetry of 1 kV has been introduced to aid visualisation. The dashed purple line is used to approximate the 0 V plane position. In the region between the guard and centre regions this plane can be seen to bend.

at the highest possible voltage, to maximise the molecular polarisation. Given these conflicting constraints, the solution was to run the centre of the interferometer at high voltage, performing rf transitions at a low voltage. Additionally, the guard regions were operated at different electric fields. The subsequent use of different frequency rf transitions was used to destroy Ramsey type $|0\rangle \leftrightarrow |1, \pm 1\rangle$ coherence.

Order of magnitude estimate

To estimate the size of this effect, the interferometer phase must be expressed in terms of spatially varying fields. This is because as the molecules propagate through space they will tend to remain aligned to the electric field $\vec{E}(y[t])$. In doing so, the molecules sample a different projection of the magnetic field, which given the small scale of the EDM interaction, will likely dominate the interferometer output. This may occur inside the electric field plates, due to the finite precision machining of the field plates, or patch potentials, but is likely to be a much stronger effect in the region between the guard and centre field plates. Given this, here I introduce a simplified model to make the analysis more intuitive using the approximation that the molecules are either aligned to \vec{z} (well inside the plates), or they are rotated some angle θ (in between the plates).

Using the known the plate geometry and voltage tolerances we can estimate the angle of rotation θ between the plates. As can be seen from figure 4.1, the required configuration is that the c and g region plates have an overall the 0 V plane offset. To lead to a systematic, this must reverse, or at least change orientation with E , as this leads to θ having a dependance upon the state of the electric field. Assume four applied voltages, $\{V_{c+}, V_{c-}, V_{g+}, V_{g-}\}$, let $V_{c\pm} = V_{g\pm} = \pm 2$ kV. In the guard region let V_{g+} acquire a small positive offset, becoming $V_{g+} + \delta$. Likewise add an offset to the negative potential such that $V_{g-} + \delta$. This places the 0 V plane in the guard region slightly closer to the negative plate. The plates themselves are identified as N and S . Upon reversal of the field, the voltages stay the same, but are now applied to the opposite plate in each region. This means that the 0 V plane in the guard region moves towards the new negative plate. The result of this process is that between the guard and centre regions, the molecules rotates opposite angles, $\pm\theta$, depending on the state of the electric field.

At running voltages of $V_g = V_c = 2$ kV, and $\delta = 1$ V, the worst asymmetry is described in table 4.1 below.

E state	true		false	
Region	g	c	g	c
Plate N (kV)	2.001	2.000	-1.999	-2.000
Plate S (kV)	-1.999	-2.000	2.001	2.000

Table 4.1: Field plate voltages for asymmetric electric field systematic. Under these applied voltages, the curvature of the electric field would lead to a rotation of the molecules upon field reversal.

To calculate the rotation of the molecule, we take the plate spacing to be d and the length of the rotated segment to be l . We can calculate θ as:

$$\theta = \arctan \left[\frac{\delta V \cdot d}{2V_{(g,c)} l} \right]. \quad (4.4)$$

For the values given in table 4.1, $\theta = 0.008^\circ$. Given this rotation angle, an interesting quantity to

calculate is the magnitude of *static* B_y required to induce an EDM of order $d_e = 10^{-28}$ e.cm.

Taking the EDM to be zero, such that $E.B$ is due solely to the Zeeman term and $B_z = 0$ T we know then that any phase measured is solely due to the projection of B_y onto the rotated axis of quantisation. Although the interferometer is 70 cm long, we only need consider the short section where the field is rotated.

Let us solve for the maximum permissible B_y in terms of d_e :

$$d_e \eta E_{\text{eff}} \frac{t_{\text{total}}}{\hbar} = \mu_B B_y \sin \theta \frac{t_l}{\hbar}, \quad (4.5)$$

$$B_y = \frac{t_{\text{total}} \eta d_e E_{\text{eff}}}{t_l \mu_B \sin \theta}. \quad (4.6)$$

Substituting in typical experimental values above, we find that the maximum permissible internal magnetic field, B_y is 7.0 nT.

It is clear that if the experiment is engineered perfectly, under homogeneous B_y , the interferometer picks up a cancelling phase at the second c - g gap, so the overall systematic EDM is zero. However, if the magnetic field B_y is different at each gap in the field plates, the cancellation is partial. In this case, for two cancelling regions, we require a difference in $\Delta B_y = 7.0$ nT between each gap in the field plates.

We might optimistically consider the magnetic field shielding to be a factor of 1000, but B_y shielding is likely to be much less due the lack of end caps on the outer shield, and access holes. If we assume a shielding factor of 10, then the maximum external field difference is $B_y=700$ nT. This is the equivalent of a 20 cm diameter loop, carrying 3.5 A, placed approximately 30 cm below the first $g - c$ region. This magnitude of magnetic field could plausibly arise from nearby turbo pumps, as well as residual fields after degaussing, so we ought to be concerned.

To test the theory, we deliberately applied a magnetic field gradient⁶ and suitable electric fields. A 20 cm, 15 Amp.turn coil was placed approximately 30 cm below the first $g - c$ field transition region. Again, assuming a shielding factor of 10, the field B_y is 300 nT at the first region of curved field, but negligible at the second $c - g$ region, giving us a suitable field difference. We then acquired a small amount of EDM data in various field configurations, with $\delta = 1$ kV, $V_g = 2$ kV, $V_c = 8$ kV. Due to the field control instrumentation used, It was only possible to add an asymmetry to the guard (g) region plates. The voltages used are shown in table 4.2 below.

E state	true		false	
Region	g	c	g	c
Plate + (kV)	3	8	-2	-8
Plate - (kV)	-2	-8	3	8

Table 4.2: To test the asymmetric electric field systematic model, an experiment was undertaken with deliberately large asymmetric potentials.

⁶This data was taken on the 14th June 2006.

The results of the experiment are shown in figure 4.2. The error bars on the measured ‘EDM’ here are weighted variances derived over each dataset. From equation 4.6 above we might expect to see an EDM of order 10^{-25} e.cm. Furthermore, changing the sign of the applied magnetic field ($B_y \rightarrow -B_y$) should change the sign of the induced EDM. It is clear from this figure that there is an induced EDM when there is both an E field asymmetry and a B_y field gradient, so this experiment seems in good agreement with theory.

With a voltage asymmetry of 1 V we might well operate with the asymmetric potentials described in table 4.2, and may expect measure a systematic of order $d_e = 10^{-28}$ e.cm.




Figure 4.2: Experimental evidence that field asymmetries, along with a magnetic field B_y can lead to a systematic EDM. The lines added are not theory, but fits to aid grouping of data points.

Solution

The obvious way to suppress this systematic is to remove B_y . However, this is very difficult as the construction of the machine prohibits effective shielding along the y -axis. Additionally, we have no good way to measure B_y ; the molecules sample B_z , and the use of a flux gate magnetometer necessarily involves removing shielding.

The solution was to engineer new high precision electric field plates, as described in §2.2.2. By changing from the three region design to a single region electric field we suppress this effect. With the new setup it is quite hard to envisage a scenario where this type of systematic may reemerge. Any machining imperfection is necessarily static, so though the 0 V plane may not be perfectly flat, any deformations that rotate the molecules are the same in each state of E . Finally we might consider patch potentials⁷. Despite gold coating our field plates, we cannot be certain that there are no localised field inhomogeneities (e.g. patch potentials) which would serve to rotate the molecules similarly to the three region electric field systematic. The best approach to fully investigate these would be to measure the electric field throughout the machine. This, along with optimisation of rf transitions motivates the development of the field mapping technique presented in the following chapter 5.

⁷Patch potentials arise on the surface of metal plates due to the granular structure of the metal surface.

4.2.2 Electric field asymmetry systematic

In Hudson's([16] p.87) original EDM experiment he noted a systematic effect in which if the interference lineshape contrast depends upon the E-field direction, and there is a non-zero B-shift, a fake EDM is observed. There is no suggestion that the interferometer input contains an additional phase term which mimics the EDM phase, rather that our simple 8 point measurement scheme cannot accurately resolve all types of shifts and scalings. In the EDM data presented in chapter 6 we correct for this effect, and so I repeat the explanation of the systematic here.




Figure 4.3: *a*: An E shift leading to a change in fringe contrast with E state, this can be caused due to a change in electric field, and hence operating slightly off resonance in one state. *b*: Introducing a B-shift, due to poor tuning of the magnetic field offset. *c*: Combined these effects lead to an apparent EDM - a change in intensity correlated with E and B.

The systematic is best understood by considering the interferometer lineshape. Figures 4.3. *a, b* depict two common situations. For simplicity the *DB*-step is not considered. In the first case the interference curve ‘changes’ with the electric field state. This is likely due to imperfect field reversals leading to the interferometer being closer to resonance in one state than the other, such that the fringe contrast is state dependent. The second case is non-zero *B*-shift: the fringes are not centred, so there a different intensity is measured on each side of the fringe. If these two effects both occur, then we find that there is a fake EDM signal. Labelling each point according to the state of the fields, we see the EDM analysis $\phi_{EDM} = B^+E^+ - B^-E^+ - B^+E^- + B^-E^- \neq 0$, despite the fact we have not introduced a real EDM phase term.

All is not lost. Terming the maximum and minimum interferometer outputs $I(\phi = 0) = I_{\max}$, $I(\phi = \pi/2) = I_{\min}$, the lineshape intensity is defined $\gamma = (I_{\max} - I_{\min})/I_{\max}$. The false EDM is then calculated as:

$$EDM_{\text{false}} = \frac{1 - \gamma}{1 + \gamma} B_{\text{shift}}, \quad (4.7)$$

which holds for any lineshape which scales and shifts uniformly. Additionally, γ can easily calculated

$$\frac{1 - \gamma}{1 + \gamma} = \frac{E.CAL}{CAL}, \quad (4.8)$$

where *E.CAL* is a compound analysis channel⁸.

To suppress this systematic, EDM_{false} is calculated for each block and subtracted from the EDM analysis. All EDM data published is corrected with respect to this systematic effect.

4.2.3 Investigation of analysis channels across the time of flight

Ordinarily, the EDM analysis integrates over the TOF. We assume that each molecule in the pulse undergoes an identical experiment, that the rf transitions are homogeneous, and each molecule integrates the same E and B fields throughout the interferometer. We therefore also assume that the value of any measured EDM is consistent across the pulse. If the various analysis channels have any structure across the TOF it is possible that they might mask a systematic EDM (perhaps the first half of the TOF has a huge positive EDM, the second a large negative EDM). This situation is clearly as unsatisfactory as having a more obvious systematic.

As part of our systematic checks a more detailed analysis of the TOF was performed. For each shot fired, the PMT records a time varying signal. The fact that the TOF has a width implies that throughout the experiment the molecular pulse has a spatial distribution. Here it is assumed that the molecules were created instantaneously from a point source, such that the pulse width is entirely due to the distribution of velocities arising from the translational temperature of the pulse⁹. Different molecular velocities imply that during the rf pulse, the molecular pulse is spread out a little (the molecular pulse length is approximately 3 cm at first transition, 8 cm at the second). Therefore the interferometer lengths and positions associated with the leading and trailing edges of the pulse are slightly different. If B , E and the rf fields are inhomogeneous, we might expect to see different interferometer phases accrued by different velocity classes, which sample different portions of the machine.

The investigation takes several hundred blocks drawn from large EDM data runs. Here I compare two sets: *set A* (2006, rf loops and three region E field), and *set B* (2007, rf transmission line and single region field). As well as differing hardware implementations, the experimental approach in the second, more recent data set was to move the first rf pulse away from the source into a region of more homogeneous magnetic field.

To extend the analysis to look for variation of signals across the TOF is relatively straightforward, though computationally expensive. Once a list of blocks to be included in the analysis has been generated, the analysis gates are set to a small portion at the edge of the TOF. The dataset is then repeatedly analysed, incrementing non-overlapping gate positions through the pulse. This yields the usual diagnostic channels, though now at various times across the pulse.

The total signal (sum of all points), will obviously vary across the pulse, due to the approxi-

⁸The *E.CAL* channel is the analysis correlated to both the E and DB steps. It represents a change in fringe gradient which coincides with a change in state electric field state.

⁹Other factors contribute to the width of the recorded time of flight including probe beam width and where the molecules are created. This is discussed in section 5.1.2.

mately gaussian form of the TOF profile. This in turn suggests that all other channels should be normalised. Consider the *DB* step, the gradient of the fringes. Ignoring noise, this is the difference between two count rates. Towards the wings of the pulse, these rates are going to be scaled by the pulse envelope. The analysis of these channels (using the rf transmission line data set *B*) is shown in figure 4.4. It is clear in these figures that the total signal and fringe gradient (*DB*) vary across the pulse, though the normalised quantity is more homogeneous.




Figure 4.4: Time binned analysis of *SIG* and *DB* channels, using loops (dataset *A*). As can be seen, the signal decays towards the edges of the pulse, as does the *DB* analysis. The normalised signal is more uniform.

In this data we can clearly see that both the signal, (*SIG*), and *DB* diminish away from the centre of the pulse due to the profile of the molecular pulse. Error bars are shown on this figure, though they are not visible on the scale of the plots¹⁰.

Our most important analysis is the *E.B* PMT channel. A non-zero result in this channel implies a non-zero EDM. Using the same EDM data sets as before the *E.B* PMT channel was calculated across the pulse (normalised against *DB*). The results are shown in figure 4.5. This data supports that there are no significant features across the pulse; there does not appear to be a velocity dependant EDM signal¹¹.




Figure 4.5: Comparative time binned analysis of EDM (*B.E/DB*) channel using both data sets. For both sets of data, the analysis is consistent with a zero EDM across the pulse. In units of *e.cm*, the confidence interval on each data point is approximately $\pm 10^{-27} e.cm$. Data points have been joined in this figure not to suggest an interpolation, but to aid grouping of the two data sets.

¹⁰For a bin drawn from the centre of the pulse, the *SIG* analysis is $20 \pm 0.025 V \mu s$.

¹¹It should be noted that for a given velocity class sampled here, the confidence interval is of order $\pm 10^{-27} e.cm$, so a smaller systematic effect could potentially still exist.

Time dependant analysis of PMT B channel

An interesting analysis is that of the PMT B -channel. Although we are primarily concerned with the $B.E$ channel, systematics can arise from the coupling of B - and E - channel effects, similar to the previous multiple region field electric field systematic (§4.2.1).

Assuming an idealised experiment the B channel represents a change in signal due to a magnetic field which shifts the fringe centre by some phase - a so called B -shift. To convert the B - channel (measured in $V\mu s$) to a magnetic field the B channel must be normalised against CAL , the gradient of the fringes.

$$B_z = \frac{B}{DB} \frac{100 \pi}{560} \frac{\hbar}{2\mu_B t}, \quad (4.9)$$

where B_z is the measured magnetic field, B and DB are analysis channel results. The term $100\pi/(4 \times 560)$ is the phase angle of the DB step.

In the figure 4.6 the two data sets described above are compared. Figure 4.6.a is the raw B channel analysis of data taken with both the transmission line and rf loops. This signal is converted to magnetic field units in figure 4.6.b. Inspecting the normalised analysis, the transmission line data shows approximately 1 nT variation across the pulse. By comparison the loops have a similar gradient, though the usable spatial extent molecular pulse is less. Error bars are presented here, though are barely visible on this scale.

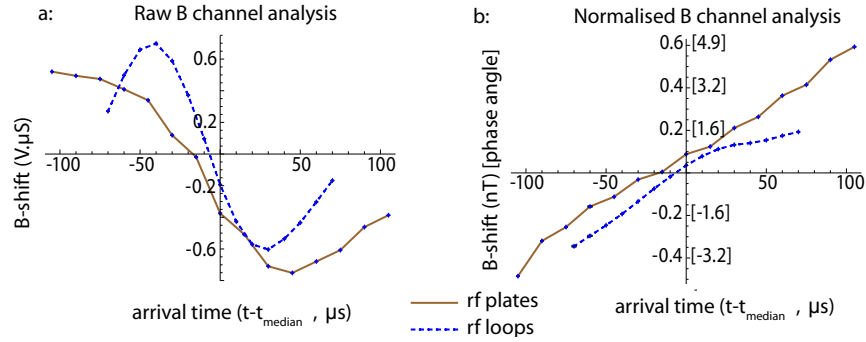


Figure 4.6: Comparative time binned analysis of the B - channel using rf plates and loops. *a*: Raw B channel analysis. *b*: B channel analysis normalised to the DB - channel. The bracketed axis labels are the raw phase in degrees. The small oscillations apparent on the rf plate data in figure *b* are artifacts of the analysis. The bin widths chosen were not a multiple of the sampling period ($1/[\text{sample rate}]$), which lead to different numbers of samples falling into adjacent bins.

There are two possible causes for the variation of apparent B -shift across the pulse. The first is that molecules from different velocity classes sample different regions of magnetic field. The presence of a field gradient in B_z along the y -axis would lead to features similar to those seen here. Secondly the interferometer can accrue a phase term due to the relative orientation of each rf field. This would appear identical to a B -shift.

Effects associated with curved rf fields A B -shift type effect can arise due to the curvature of the rf magnetic field. We assume initially that the rf loop is a perfect circular loop, orientated in the $x - z$ plane. The molecular beam passes through the centre of the loop. Away from the plane of the rf loops, the magnetic field component of the rf field curves away from the y -axis. This is shown in figure 4.9, where a green ellipse represents the approximate spatial extent of the molecular pulse at the first transition. As the pulse moves away from the centre of the x - axis molecules experience a curved rf field.




Figure 4.7: Magnetic field surrounding an rf loop. Off axis it can be seen that the magnetic field vector has a radial component. The green ellipse represents the extent of the molecular pulse at the first loop.

In a two pulse rf interferometer scheme, a relative rotation of the rf field is manifest by an additional interferometer phase term. The concern is that for off axis molecules, the applied rf field is rotated, such that the two rf transitions are not parallel. We need to calculate the effect of misaligned rf pulses.

To illustrate this, the Hamiltonian describing the interaction between an arbitrarily orientated magnetic field and the molecule was derived. The interaction between the molecules, and an external field is given by

$$(\mu_B \vec{S} + \mu_I \vec{I}) \cdot \vec{B}(t). \quad (4.10)$$

As $\mu_B \gg \mu_I$, the nuclear spin is ignored. I label the basis $|F, m_F\rangle = (|0, 0\rangle, |1, 1\rangle, |1, -1\rangle)^T$, (the state $|1, 0\rangle$ plays no part in the experiment so is ignored). Using this Hamiltonian, a π -pulse (Π) is defined as:

$$\Pi = \frac{1}{2\sqrt{2}} \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} B_x(t) + \frac{1}{2i\sqrt{2}} \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} B_y(t) \quad (4.11)$$

Here as I have ignore the $|1, 0\rangle$ state, and hence do not include the B_z field interaction. A pure

B_x driving field will generate a superposition of states $1/\sqrt{2}(|1, 1\rangle + |1, -1\rangle)$, with the application of the full interferometer sequence leading to the lineshape described in chapter 1. If the axis of quantisation is rotated a small angle θ to the rf field, the component of B_y produces a state

$$\Pi|0, 0\rangle = \frac{1}{\sqrt{2}} (e^{i\theta}|1, 1\rangle + e^{-i\theta}|1, -1\rangle). \quad (4.12)$$

The application of an rf field at angle θ to the axis of quantisation creates a superposition with an initial phase θ . This is distinct from the phase ϕ which arises from the evolution of the superposition throughout the interferometer. As with the EDM experiment, between the rf pulses, the states precess with an evolution operator U . A second pure B_x π -pulse now allows us to measure the total interferometer phase, which is now $\theta + \phi$.

$$\Pi U \Pi|0, 0\rangle = \frac{1}{\sqrt{2}} (\cos(\theta + \phi)|0, 0\rangle + \sin(\theta + \phi)|1, \pm 1\rangle). \quad (4.13)$$

This is similar to the usual EDM interferometer output described in equation 1.3.e, but with an additional phase term θ . In order to relate this to the B -channel variation across the pulse, we need to consider an rf field which is rotated differently for each velocity class. The obvious candidate is the rf field off-axis, as shown in figure 4.9.

Figure 4.6 shows the B -channel analysis correctly signed for the manual reversal of the magnetic field¹². Thus figure 4.6.a represents the contribution of a magnetic field gradient. However, for systematic effects that appear as a B -shift, but are not dependant a ‘real’ magnetic field, we must not sign the analysis.




Figure 4.8: Block level B -channel analysis unsigned by B -state. This low level inspection of the B -shift shows that each block is clearly effected by this systematic. The blocks are grouped into those where the manual B -reversal is ‘True’ (orange), and those where it is ‘False’ (blue). The thick curves show the respective weighted means of each set.

In figure 4.8, the unsigned analysis of the B -shift of each block is presented. As can be seen the analysis falls into two groups, which correspond to the state of the manual B reversal. The thicker curves represent the weighted mean of each data set. Taking the means of these curves we can generate a similar figure to 4.6.

¹²Recall both magnetic and electric field polarities are periodically reversed (§2.2.1). This B reversal reverses the sign of each phase that the interferometer is measured at. To correctly sign the B -shift of a block it must be signed

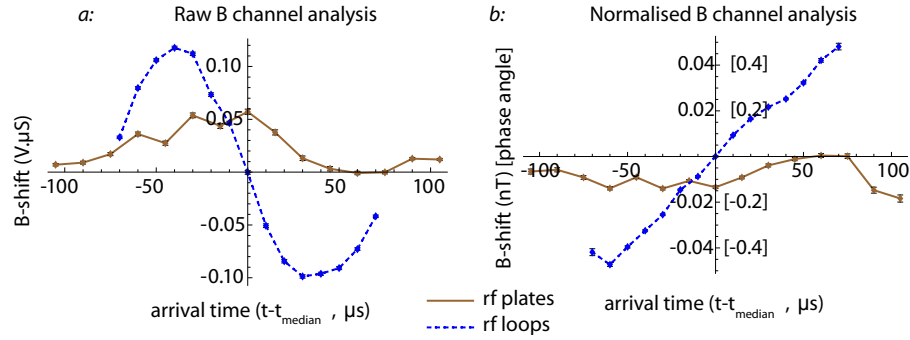


Figure 4.9: B-channel analysis unsigned by B-reversal state. In comparison to the signed analysis (figure 4.6, which probe real magnetic fields), it is clear in this figure that the transmission line rf data has much less structure than the rf loop data. In figure *b*;, the loop varies approximately linearly across the pulse, though, compared with the unsigned data, with approximately one tenth the gradient.

The unsigned analysis is clearly very different to the analysis presented in figure 4.6. The transmission line data and loop data sets no longer look similar. The transmission line data has no strong structure across the pulse, whereas the loop data retains an obvious gradient. Given that similar quantities of data were taken in each state, it is reasonable to conclude that the rf loop data here suffers from a systematic unrelated to the magnetic field gradients.

If we assume that therefore the curvature of the rf field is *solely* responsible for this, we can estimate the approximate displacement of the molecular beam from the centre of the rf loops.

Taking a single current loop as a model for the rf loop antenna it is possible to calculate the rotation of the magnetic field at any point in space. A path through the centre of the loop, orthogonal to the plane of the loop, has zero rotation of the field. Off axis however there is an angular dependence of the field. Therefore, simplifying the molecular beam to a 1-D line, we can fit the observed phase shift to the angular component of the rf field, allowing us to approximate how far the molecular beam is from the centre of the rf loops. The results of this are shown in figure 4.10. To simplify the analysis, I assume that the molecular beam is perfectly aligned to the second rf loop, such that the applied rf is not curved. In this situation the rotated/curved rf field associated with the first rf transition imprints a phase across the pulse, which is then measured along with the free evolution phase at the second transition.

The investigation of the rf loop data using the binned analysis revealed a phase shift across the pulse, which we have first considered as arising from a magnetic field gradient. The interferometer phase shift associated with the magnetic field ϕ_B , ranged approximately linearly from $-0.52^\circ \leq \phi_{B\text{eff}} \leq 0.51^\circ$.

Under the approximation that the beam is a thin, one dimensional ray, the molecules experience a curved rf field when the beamline is not collinear with the axis of the rf loop. Fitting of the phase shift data, $\phi_{B\text{eff}}$ suggests that the beam is translated 0.33 mm from the centre of the either

by the state of the manual *B*-state.

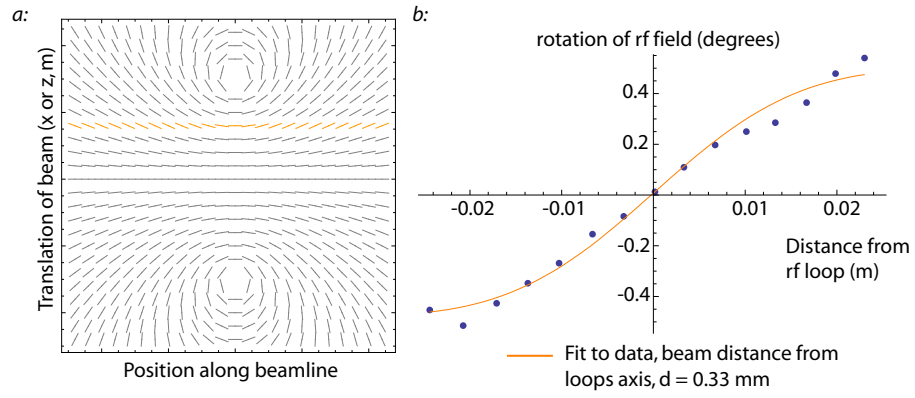


Figure 4.10: *a*: Vector plot of static magnetic field surrounding a current loop. The orange points illustrate the rf field orientation experience by an off-axis molecule. *b*: Fit of the B/DB phase data to the rf field rotation, parameterised by lateral displacement of the beam. The data suggests the lateral displacement of the molecular beam is 0.33 mm away from the centre of the loop.

rf loop (figure 4.10) in the $+x$ direction. Such a level of misalignment is entirely possible.

Though this brings some insight, the actual situation is much more complicated. The most obvious problem is that with two (identical) rf fields, the effect partially cancels. Only if one rf transition has imperfect timing should a net effect be observed. Then, we must consider the expansion of the pulse through the machine. It is impossible to estimate the rf field vector that a given velocity class experiences.

The recent move to the transmission line rf scheme has evidently suppressed this effects. The TEM mode rf is sufficiently homogeneous that the B -shift seen in figure 4.9 is certainly not related to inhomogeneous rf fields.

However, this says little of the signed B -channel analysis, which leads to features an order of magnitude larger than those discussed here, in both the transmission line and loop data sets. To investigate these we must consider the rôle of magnetic field gradients.

Magnetic field gradients Gradients in B_z along the y -axis can lead to changes in the measured B -shift across the pulse. Assuming all molecules originate from a point source, we can for a given velocity class define a region of space over which they were in the interferometer stage of the experiment, hence sensitive to magnetic fields.




Figure 4.11: Using the arrival time information and rf pulse times the location of the molecules at the start and end of the interferometer can be calculated. The solid lines represent the interferometer locations for the slow molecules, the dashed lines the fast molecules.

Shown in figure 4.11 are the interferometer positions for the leading and trailing edges of

the molecular pulse. There are two important features shown in this figure. Firstly that the transmission line and loop data sets were taken with different rf timings, and therefore are sensitive (i.e. in a coherent superposition) over different regions of the machine. Secondly, that the fastest and slowest molecules from each set (dashed and solid lines) are not sensitive over the same regions. The fast molecules integrate a little ‘extra’ at the top of the machine, the slow molecules likewise at the bottom. Clearly there may exist some form of magnetic field which could lead to the B -shifts across the pulse seen in 4.6.

To probe this effect a technique was developed for measuring the magnetic field component B_z throughout the interferometer region, where z is defined by the electric field vector. This technique is discussed fully in chapter 5. Using this field map it is possible to calculate the anticipated B -shift measured across the molecular pulse for both rf loops and plates.

$$B_{\text{shift}} = \frac{\int_a^b B_z(y) dy}{b - a}, \quad (4.14)$$

$$a = l \frac{\tau_{\text{rf1}}}{\tau_{\text{arrival}}}, \quad b = l \frac{\tau_{\text{rf2}}}{\tau_{\text{arrival}}}, \quad (4.15)$$

where τ_{rfx} is the rf pulse time, τ_{arrival} is the arrival time and l is the distance from the source to the probe PMT. Using the timing information for the rf loops ($\text{rf}_1=840 \mu\text{s}$, $\text{rf}_2=1955 \mu\text{s}$) we anticipate approximately 15 nT variation across the pulse. However, using the rf transmission line timings, ($\text{rf}_1=1020 \mu\text{s}$, $\text{rf}_2=1850 \mu\text{s}$), we expect only 1 nT variation.

The rf loop data analysis is then inconsistent. The analysis of the EDM data B -shift suggests a shift of 0.6 nT across the pulse - here we have a shift of 16 nT. The most likely explanation for this is that the magnetic field has changed over the 1-2 years between taking the EDM data. Though inspection of the EDM data sets suggest that the B -shift across the pulse is quite stable over time, it may be that in upgrading the experiment the magnetic field changed significantly, making this analysis useless.

Contrastingly, the rf transmission line analysis is more cohesive. With no susceptibility to curved rf fields, magnetic field gradients must be suspected. Using the recent magnetic field map, created just before the acquisition of the EDM data, a reasonable comparison between the EDM data B -shift and the calculated B -shift may be made. Both display a shift of approximately 1 nT across the TOF. Though not attempted here, EDM data could be taken with a range of rf timings to allow a more complete comparison with the magnetic field map.

4.3 Systematic effects conclusion

Here I have presented recent work in understanding and controlling several new systematic effects. The first systematic investigated demonstrated how asymmetric electric potentials in the three

region field scheme, combined with a static magnetic field can lead to a systematic EDM. To suppress this effect the electric field plates were redesigned. The use of a single region now removes the possibility of systematics associated with the curvature of the 0 V plane.

An artefact of our acquisition scheme which wrongly interprets certain shifts as an EDM was previously discovered by Hudson. This is presented here for completeness.

Finally a systematic effect affecting the B - channel was discovered, whereby structure was seen in the time binned analysis. Though not a systematic EDM, the effect was investigated. The most probable cause for this effect was magnetic field gradients towards the end of the machine.

Through the discussion of these systematics, the general theme has been that a greater understanding of the electromagnetic environment is beneficial. Additionally, a map, or even partial map, of the electromagnetic fields used in the experiment could form the basis of a more sophisticated model of the experiment. With this in mind, before discussing the EDM measurement, I present techniques developed for mapping of electric, magnetic and radio frequency fields throughout the experiment.

Chapter 5

Measurement of electric, magnetic and rf fields.

This chapter presents techniques for mapping electric, magnetic and radiofrequency fields throughout the interaction region¹. These field mapping techniques are described in the Hudson et al. paper [47] ‘Pulsed beams as field probes for precision measurement’, though here we also present an implementation of Ramsey’s separated oscillatory field technique for more accurately probing the average electric field throughout the interaction region.

5.1 Field mapping using a single rf resonance

A technique has been developed to map electric, magnetic and radiofrequency magnetic fields throughout the interaction region of the experiment. The principal motivation for making these measurements is to allow us to develop a realistic model of the experiment to aid us in understanding possible systematic effects.

The basic concept behind this field mapping scheme is that we can use a single rf transition to probe Stark and Zeeman effects in YbF, and therefore measure electric and magnetic fields. Additionally, we can also probe the amplitude of the rf field by measuring the required power to drive exactly a π pulse. Crucially, because we know the timing of the rf pulse, the location of the source and the velocity of the molecules, we in fact know the position of the rf transition, allowing us to measure the electromagnetic environment at a specific point in space.

5.1.1 Scanning a single rf transition

First consider a simple experiment, in which a pulse of YbF molecules passes into the interaction region, is subject to an rf pulse, and goes on to be probed using LIF. To measure the E and B

¹The interaction region is defined as the shielded region, where magnetic, electric and rf fields are well controlled - see figure 1.8.

fields, our requirements are that we can drive an rf transition between two states and that we can probe the final state populations². There is no reason not to use the states relevant to the EDM experiment, as we obviously have the correct pump/probe lasers and rf apparatus. To label the states, use the notation $|F\rangle$, or when appropriate $|F, m_F\rangle$.

The first step, as with the EDM experiment is to prepare the molecules in the $|0\rangle$ state. With the probe laser tuned to the to the $|1\rangle$ level, this state is pumped to the next electronic state, effectively removing these molecules from the experiment. Next, rf is applied to drive transition $|0\rangle \rightarrow |1, \pm 1\rangle$. finally the pulse is probed on $|1, \pm 1\rangle$ using LIF. A summary of such an experiment is shown in figure 5.1.




Figure 5.1: A simplified single resonance rf experiment. The molecular pulse enters from left (population of $|1\rangle$ state shown). This is then pumped, leaving the $|0\rangle$ state populated. A resonant rf pulse then drives this population to the $|1\rangle$ state (π -pulse). The $|1\rangle$ population is detected using LIF.

However, this so far tells us very little. In order to infer details of the fields we need in the very least to measure the lineshape centre frequency. To do a scan must be performed, where on each successive shot the rf frequency is incremented³. The data recorded from such a frequency scan is 2-dimensional: a series of TOF profiles, recorded at increasing rf frequency - see figure 5.2 below. For most purposes we are not interested in the timing information, so the integrated TOF is calculated. This is just taken as the numerical integral of the TOF signal. An example of fitting such a scan is given figure 5.3. We choose a range of rf frequencies to scan over based upon the linewidth and centre frequency. The rf power used to drive the transition must be carefully set, such that on resonance a π -flip is driven. In a typical experimental configuration, the rf transition takes 18 μs with around 30 dBm forward power to the rf plates.

Figure 5.3 shows the integrated TOF is fitted to the theoretical lineshape, equation 3.1.f, yielding the transition frequency.

Assuming a very rapid rf transition performed at time t_{rf} , we can calculate the position (y_{calc}) of the molecules at the time the rf was applied⁴.

²For mapping the magnetic field we need the need to be able to resolve the Zeeman splitting of the $F=1$ state.

³Recall, in EDM nomenclature, a single experiment is a shot; a series of shots is a scan.

⁴We are also temporarily assuming that all molecules were born simultaneously, at a single point in space, and that the LIF detection scheme has high resolution in the y direction. We will investigate the implications of these assumptions shortly.




Figure 5.2: Radiofrequency scan over the $|0\rangle \rightarrow |1, \pm 1\rangle$ transition. The surface is the raw 2D dataset. The two projections are the rf scan (pink), and the TOF (red).




Figure 5.3: Fit of single transition rf lineshape. A series of shots were recorded over a range of rf frequencies. The PMT signal from each shot was integrated over the TOF, and plotted. This data was fitted to the rf transition lineshape (equation 3.1.e), yielding the centre frequency from which we can calculate the electric field.

$$y_{\text{calc}} = \frac{l_m t_{\text{rf}}}{t_{\text{arr}}}, \quad (5.1)$$

where l_m is the distance from the source to the probe PMT.

The velocity distribution causes the pulse to expand from 3 cm to 8 cm during its flight through the interaction region, so to a first approximation a single rf scan can resolve at least 3 cm of beamline. To map subsequent sections, the rf timing is adjusted such that the pulse is applied slightly later, when the molecules have progressed further along the interaction region. This technique requires that rf transitions can be driven throughout the region to be mapped. This was initially attempted using an array of four rf loops throughout the machine. Though this worked adequately, but was laborious, with poor quality rf transitions. The new transmission line rf plates are far superior, allowing homogeneous rf transitions to be driven throughout the interaction region.

5.1.2 Mapping the electric field

Consider a specific example - mapping the electric field. The first part of the procedure is to scan out and optimise the frequency and amplitude of an rf transition well inside the electric field, at least 2-3 cm away from the end of the plates. This is to ensure that for the chosen rf parameters the scan covers the entire transition and that the rf power delivered drives a π -pulse, on resonance.

The appropriate timings for locating the rf at the ends of the electric field region are easily calculated given the machine dimensions, though typically the timings vary little day to day. Starting with rf pulse timing to drive transitions at the bottom of the interaction region (near the source), the rf lineshape is scanned. This is repeated at suitable intervals along the beam⁵. We now have scans of the rf transitions for every point along the beam trajectory.

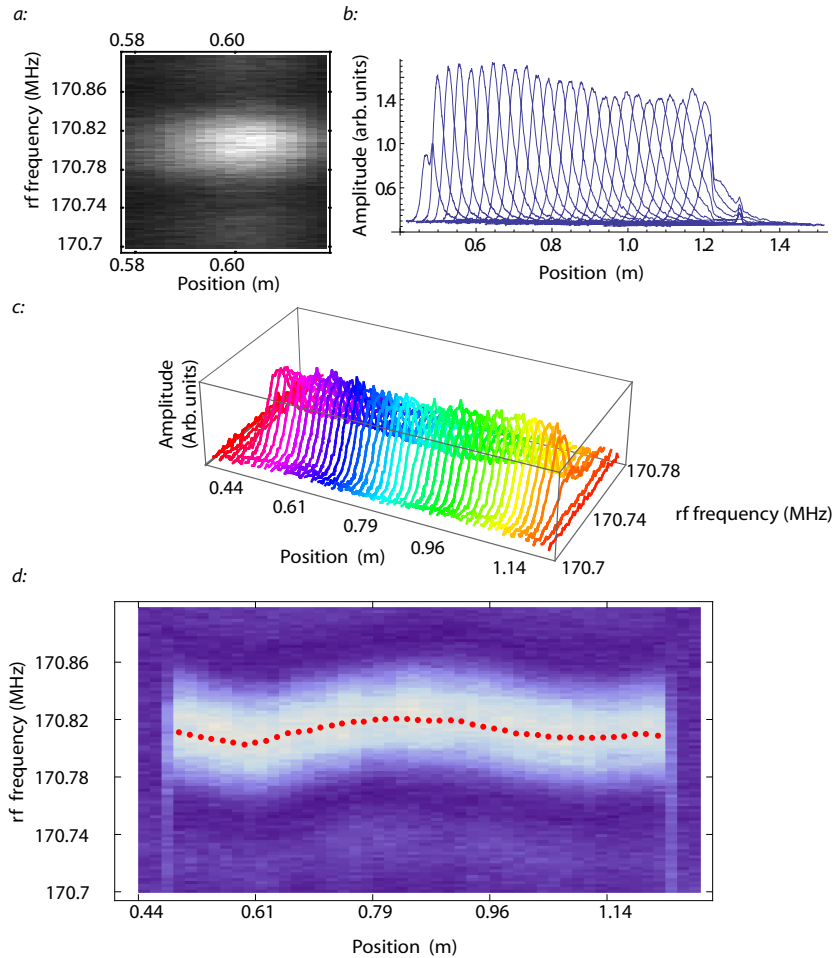


Figure 5.4: *a*: 2-D data consists of many TOF profiles, each taken at a different rf frequency
b: Integrating each scan over the frequency dimension generates the mean TOF. This can then be used as a guide to where the molecules were during the rf pulse. Clearly there is considerable overlap between scans.
c: The composite scan, appropriately binned, contains a series of independent rf lineshape scans.
d: Each bin is fitted to the rf transition lineshape (see appendix B). Here the centre frequency of the transition is marked in red. It is then trivial to infer the electric field at each position using the known Stark shift of the hyperfine transition.

⁵Because the pulse expands as it travels, we could increase the amount by which we step each scan in space, though typically, for simplicity we continue to step at 2 cm increments.

The analysis of the data is quite straightforward, and is shown in figure 5.4. We start with a single rf frequency scan. The TOF needs to be trimmed (gated), as the PMT captures a few hundred μs of noise either side of TOF (see figure 5.4.a). The timing units are then converted to position. This scan contains information about the electromagnetic environment where the molecules were when the rf was applied. Successive overlapping scans are taken throughout the interaction region. Figure 5.4.b shows the position of the pulse with successive scans of later rf timing. From this set of data we can compile the data into a single composite scan (5.4.c). The frequency scans are divided into non-overlapping spatial bins, introduced to account for the spatial resolution - we wish to ensure that the electric field measured for each position is an independent measurement (discussed below in section 5.1.2).

The plot shown in 5.4.c shows how the TOF actually contains many rf scans. Each bin is now processed independently, fitting the lineshape function to the data. From this we can determine the centre of the line, then using Stark shift calculations (appendix C.1), calculate the observed electric field associated with that bin. Repeating for each bin we generate a series of field mappings⁶ $E_z(y)$, for the electric field as a function of position.

Estimation of error bars

In the previous section, the idea of binning the data was introduced. Each rf scan contains information about a certain region of space (3-6 cm). To construct a composite we establish a series of spatial bins to simplify joining scans. An important question is how wide these bins should be.

For any point, we can assign a position uncertainty equal to half of the distance the molecules move during the π -pulse (~ 10 mm for a $18 \mu\text{s}$ pulse). Below this scale we have no spatial information. An additional smearing is produced by the source, as it is not clear that the molecules are created at the same time or point in space. To calculate the position from the arrival time use the relation $y_{\text{calc}} = \frac{l_m t_{\text{rf}}}{t_{\text{arr}}}$, where t_{arr} is the arrival time, t_{rf} is the rf pulse timing, and l_m is the machine length. Now, assume our detection time is subject to error due to the uncertainty in the origin of the molecules and their creation time, such that:

$$t_{\text{arr}} = \frac{l_m - \Delta l}{v} + \Delta t, \quad (5.2)$$

where Δt is an offset in the creation time and Δl is an offset in the creation position. Now need to calculate the difference in position between our estimated position, and their true position:

$$\Delta x_{\text{rf}} = x'_{\text{rf}} - x_{\text{rf}} = \frac{t_{\text{arr}} - t_{\text{rf}}}{t_{\text{arr}}} \frac{l_m \Delta t - t_{\text{arr}} \Delta l}{t_{\text{arr}} - \Delta t}. \quad (5.3)$$

Assuming Gaussian distributions for Δt and Δz , the standard deviation of the rf position is:

⁶The electric field defines the z -axis.




Figure 5.5: Map of the electric field with error bars; 15 mm bins. The data point nearest the source is not an error, but reflects the decay of the electric field close to the end of the plates.

$$\sigma_{x_{\text{rf}}} = \frac{t_{\text{arr}} - t_{\text{rf}}}{t_{\text{arr}}^2} \sqrt{t_{\text{arr}}^2 \sigma_l^2 - l_m^2 \sigma_t^2}. \quad (5.4)$$

We have no way of measuring the timing error σ_t for the YbF source, but for lithium hydride using similar apparatus, it was found the molecules are produced over $\sigma_t = 5 \mu\text{s}$ [48]. Likewise, if it is assumed that the molecules are produced with no spread in time, we find an upper bound $\sigma_l = 11 \text{ mm}$. We might also consider the instrument response of the PMT, as it averages over some short length of the pulse, corresponding to the width of the probe beam. This will add only a few millimetres uncertainty, so is only a small effect. We take an upper bound on the net uncertainty of $\sigma_x = 15 \text{ mm}$.

The error in the field can be derived from the asymptotic standard error of the fit for each bin. The *Mathematica* function *NonLinearRegress* was used for fitting. This approach at estimating error bars was combined with suitable width bins to enforce independent measurements, yielding the analysis presented in figure 5.5.

Electric field mapping summary

The electric field mapping implied a field of mean value 3270.3 Vcm^{-1} , with large features ranging from 3242.4 Vcm^{-1} to 3302.63 Vcm^{-1} . The applied voltage was $\pm 2000 \text{ kV}$. After smoothing small features, the maximum field gradient was found to be 3.5 Vcm^{-2} , though assuming the small features are not noise, but are in fact real defects, a gradient of up to 6 Vcm^{-2} is seen inside the interaction region. The mean error bar was found to be 1.09 Vcm^{-1} . Converting the electric field to plate spacing, implies that the plates have a ‘bowing’ of $\pm 100 \pm 4 \mu\text{m}$ (given a field sensitivity of 1.09 Vcm^{-1} , we can resolve the plate spacing to $4 \mu\text{m}$). No measurements of the electric field reversal precision have been made. However, integrating across our map we can measure the mean field with a precision 0.6 Vcm^{-1} .

5.1.3 Mapping the magnetic field

The magnetic field can be mapped in a similar manner to the electric field, using the Zeeman effect to probe $B_z(y)$. Our approach is to use the B_z coils used in the EDM experiment to generate a field throughout the interaction region. Then we use the Zeeman splitting of the $|1\rangle$ state to measure the applied field at some point in space. As before, we increment along the machine, scanning a series of rf transitions. A moderate magnetic field of 1 μT gives a splitting of the $|1, -1\rangle$ and $|1, 1\rangle$ states of 28 kHz. The transition linewidth for an 18 μs pulse is approximately 50 kHz, which is too broad to resolve this magnetic field. Therefore, a longer pulse of 90 μs is used, with a narrower linewidth of 20 kHz, at the expense of spatial resolution.

As before, we start with an rf scan well inside the machine, though now subject to the applied magnetic field. It is important to check that the splitting is resolvable, and that the entire lineshape is within the scan range. In figure 5.6, the data has been fitted to a double lineshape, the linear sum of two rf lineshapes (3.1) separated by a detuning $2\Delta\nu$.




Figure 5.6: Zeeman split rf transition. With a 90 μs rf pulse, the linewidth is 20 kHz, which allows the Zeeman splitting to be resolved. The splitting fit implies a splitting of 24.5 kHz, corresponding to a magnetic field of magnitude $|B_z| = 1.75 \mu\text{T}$.

As with mapping the electric field, a series of scans must be compiled to a single composite map. In figure 5.7 the binned rf map has been processed to fit the line centre (Stark shifted), and the splitting of the $m_F \pm 1$ sub-levels. The electric field is applied in order to align quantisation axis to the \hat{z} -axis, such that we measure the magnetic field which the molecules sample within the interferometer.

In order to properly interpret this data, the map was again divided into non-overlapping, independent bins. These are wider than for the electric field mapping, as the increased rf pulse length reduces the spatial resolution.

Finally, to convert the splitting to a magnetic field, the frequency of the splitting (δ) due to the Zeeman interaction is converted to a magnetic field.




Figure 5.7: The individual rf scans can be composed into a map of the interaction region. The red points represent the centre of the fitted lineshape, whilst the green and orange points are generated by taking the centre frequency, then adding/subtracting fitted Zeeman splitting frequency.




Figure 5.8: Magnetic field map showing the magnetic field through the interaction region.

$$\begin{aligned}
 |h\delta| &= |\mu_B B_z|, \\
 |B_z| &= \frac{h|\delta|}{\mu_B}.
 \end{aligned}
 \tag{5.5}$$

The resulting data from the magnetic field map is shown in figure 5.8. The limitation of this data is that it only describes the magnetic field inside the machine at a particular applied field. We have no way to resolve the ambient and applied fields.

Resolving between ambient and applied magnetic fields

The technique presented above demonstrates how the molecules can be used to probe magnetic fields. As described it is not possible to distinguish between ambient and applied field homogeneities, or the sign of the measured field. There is however a simple solution. After mapping $B_z(y)$, the current to the coil is reversed and the process is repeated. Assuming the ambient




Figure 5.9: *a*: Four independent maps of $B_x(y)$ were taken with B_z - coil currents of ± 60 , and ± 110 mA. *b*: From this the ambient and applied magnetic fields were extracted.

field is unchanged, with the applied field reversed we can now extract the two components of the field, where the sum of the two maps represents the ambient field, and the difference is sensitive to inhomogeneities in the applied field. To further enhance this technique 4 currents are used, $\{60, -60, -110, 110\}$ μA .

Magnetic field summary

The technique I have presented was used to measure ambient and applied magnetic fields, $B_z(y)$. The uncertainty in the ambient field was approximately 5 nT, whilst for the applied field the uncertainty was 10 pT/mA. These uncertainties reflect the errors associated with fitting, as well as experimental noise.

The results show that the applied field decays slightly towards each end of the interaction region, as the molecular pulse approaches the end caps of the μ -metal shield. Interestingly we also see unexpected features on the ambient field, which increases sharply near the source, whilst it decreases near the detector. This is due to poor magnetic shielding towards the ends of the machine. This data allows us to immediately quantify a region of homogeneous field. The spatial error bars can be derived in a similar manner to electric field, though with longer rf pulses the uncertainty increases to 40 mm.

5.1.4 Mapping the rf field

The final field that we are interested in mapping is the rf field, used to generate the π pulses required for our experiment. Our technique probes the modulus squared amplitude of the oscillating magnetic component in the x and y directions, as a function of position along the beam. To do this the dynamic evolution of the transition is monitored. In much the same way that the rf




Figure 5.10: Scan of rf amplitude for a position well inside the interaction region, using the transmission line rf system. Using LIF the population of the 1 state is probed. The amplitude of the rf was scanned logarithmical - dBm. Taking the antilog of amplitude leads to bunching of the data points seen here.




Figure 5.11: (a) Map of relative rf amplitude throughout the interaction region; fit to 170.82 MHz standing wave; (b) Single loop rf amplitude; fit to single loop DC magnetic field. Note these figures are have very different x -axis scales.

frequency scan was previously used, we now scan the rf amplitude, with the frequency of the rf set to resonance. Scanning the rf amplitude drives some fraction of a π -pulse between $|0\rangle$ and $|1, \pm 1\rangle$ (see figure 5.10).

As with the electric and magnetic field maps, this TOF data was binned, and a map of the rf field throughout the machine was then constructed. Each bin was fitted to a function of form $\gamma + \beta \sin^2(\alpha \cdot A_{\text{rf}} + \phi)$, where the scaling parameter α is used to describe the relative rf power the beamline, and A_{rf} is the forward rf power to the plates. The results are shown in figure 5.11.

As seen in figure 5.11, the rf field from the transmission line rf, though not perfectly flat, is significantly more homogeneous than the loop implementation, which rapidly decays over a few cm.

The transmission line inhomogeneity stems, we believe, from reflected rf at the end of the guide. The inhomogeneity here suggests 20% of the rf power is reflected from the end of the plates, due to the impedance mismatch as we couple the $34 \, \Omega$ transmission line into the $50 \, \Omega$ coax feed (the theoretical reflection coefficient of this mismatch is 19.0%).




Figure 5.12: The effect of a 7 kHz detuning on the mapping of the rf field amplitude. *a*: With 18 μs pulses we see that the detuning leads to a small reduction in the Rabi flopping rate. *b*: With 90 μs pulses the Rabi rate increases significantly, leading to a large error.

Radiofrequency map error bars

The position error bars and binning conditions are treated exactly as for the electric field, setting the bin widths to 15 mm. Either confidence intervals or the asymptotic standard error from the fit can be used to establish an error bar for the relative amplitude.

There is a systematic error that can affect this experiment. If at some position in space the electric field changes, we will fall off resonance due to Stark detuning, causing the Rabi rate to increase and contrast to decrease. Without detailed electric field diagnostics, an increasing Rabi frequency due to detuning is easy to confuse with an increased rf field on resonance.

Taking the derived formula for the rf lineshape (3.1) it is straightforward to calculate the pulse length required to achieve a π pulse of resonance, for a given rf power. This is then compared to the rf power to reach the peak of the Rabi oscillation curve when 7 kHz off resonance⁷. The difference in power required between these two cases is the source of our systematic. In the case of the 18 μs rf pulse, this appears as a 0.54 μs reduction in the required pulse length (a relative rf power reduction of 3.13%). With a longer 90 μs rf pulse, this appears as a 34 μs reduction in pulse length; 37.9% power reduction, a huge effect. Rabi flopping curves for these situations are shown below in figure 5.12. The two figures compare the systematic with (a) 18 μs and (b) 90 μs rf pulses. Rabi flopping curves are shown in each case for the on resonance, and 7 kHz detuned scenarios. In 5.12.a, the effect is slight. In 5.12.b a dramatic shift in apparent rf power would be observed. Though not attempted here it would be straightforward to correct for this effect by suitable detuning of the rf frequency where appropriate. This analysis suggests that when mapping the rf field a shorter rf pulse is more appropriate.

It is not clear how rf amplitude variations could lead to an EDM systematic. However, with this technique we have a good understanding of the variation of the rf field amplitude throughout the machine. This information may well prove valuable at a later stage when performing a more exhaustive search for higher order systematic effects.

⁷For a field of 3.30 kV, we expect a maximum Stark detuning due to field inhomogeneities of up to 7 kHz.

5.2 Probing the electric field using a Ramsey interferometer

Thus far we have considered mapping the electric field at a series of points in space. This is principally of use in mapping spatial variations and field gradients as discussed in chapter 4, though we can also infer a mean field to a high level of accuracy. However, we might also consider measuring the quantity $\int_p E_z(y)dy$, where p is the path of the molecules through the interferometer. This is of interest because the interferometer EDM phase component is accrued along the length of the interaction region. We might also like to measure this quantity in each electric field state to investigate reversal symmetries.

Our approach here has been to implement a double resonance separated oscillatory field experiment⁸, not dissimilar to how we measure the EDM. The EDM experiment uses a so called π -pulse to drive the transition $|0\rangle \rightarrow \frac{1}{\sqrt{2}}[|1, -1\rangle + |1, 1\rangle]$. This superposition of hyperfine states is sensitive to splitting of the hyperfine magnetic sub-levels, not the global Stark shift. However, if we drive a $\pi/2$ -pulse, we form the superposition $|0\rangle \rightarrow \frac{1}{\sqrt{2}}[|0\rangle + |1\rangle]$. The energy difference between these states *is* dependent on the electric field, and as such one would expect that the states of the superposition develops a relative phase

$$|\psi\rangle \rightarrow \frac{1}{\sqrt{2}}[|0\rangle e^{i\frac{\Delta\omega_E t}{2\hbar}} + |1\rangle e^{-i\frac{\Delta\omega_E t}{2\hbar}}], \quad (5.6)$$

where $\Delta\omega_E$ is the precession frequency due difference in energy between the $F = 0$ and $F = 1$ states. Through careful use of rf transitions, we can probe this phase. This idea is developed in the next section.

5.2.1 Dynamics of the separated oscillatory field approach

The molecule can be treated as a two level system, $F=0$, $F=1$, allowing us to use Ramsey's analysis directly [49]. We could use the rf transitions previously derived for the 4-level system which fully describe the amplitudes of the $|F = 1, m_F = \pm 1\rangle$. However, any splitting of the magnetic sublevels has no net effect on the output of the Ramsey interferometer, and hence a two level analysis is sufficient.

Throughout this section I will use the notation $|0\rangle$ and $|1\rangle$ to represent the basis $|F = 0\rangle$ and $|F = 1\rangle$.

As with the EDM experiment we first prepare our molecules using a pump beam to remove all molecules in the $|1\rangle$ state. We use the same light source to probe the $|1\rangle$ state using LIF.

The rf pulse is described in equation 5.7. Here we assume that the rf frequency is resonant with the $|0\rangle \leftrightarrow |1\rangle$ transition frequency, $\omega_{\text{rf}} = \omega_0 = (E_1 - E_0)/\hbar$. In order to achieve exactly a $\pi/2$ -pulse, the Rabi interaction term b is set to $b = \frac{\pi}{4\tau}$, where τ is the rf pulse length and t is the

⁸This is also known as a Ramsey interferometer.

time the rf pulse is applied.

$$R_{\frac{\pi}{2}}(t, \tau) = \frac{1}{\sqrt{2}} \begin{pmatrix} e^{\frac{1}{2}i\tau\omega_{\text{rf}}} & ie^{\frac{1}{2}i(2t+\tau)\omega_{\text{rf}}} \\ ie^{-\frac{1}{2}i(2t+\tau)\omega_{\text{rf}}} & e^{-\frac{1}{2}i\tau\omega_{\text{rf}}} \end{pmatrix}. \quad (5.7)$$

The experiment begins in the state $|0\rangle = (1, 0)^T$. The first rf pulse creates the superposition:

$$R_{\frac{\pi}{2}}(t, \tau) \cdot (1, 0)^T = \frac{1}{\sqrt{2}} \begin{pmatrix} e^{\frac{1}{2}i\tau\omega_{\text{rf}}} \\ ie^{-\frac{1}{2}i\tau\omega_{\text{rf}}} \end{pmatrix}. \quad (5.8)$$

The next step is to allow this superposition to evolve as the molecules propagate through the interaction region. This operation can be trivially calculated by taking the matrix describing the interaction of a dipole with an oscillating rf field, where the interaction b is zero. In sections 5.1.2, the variations of the electric fields inside the interaction was measured. It is clear then that the molecules will, as the propagate through the interaction region, experience varying electric fields. As a result the instantaneous precession frequency will vary. However, as we only observe the accrued phase at one point, the detection region, we can only infer the mean electric field, corresponding to the angular frequency $\omega_{\Delta\bar{E}_z}$.

$$U(t) = \begin{pmatrix} e^{\frac{1}{2}it\omega_{\Delta\bar{E}_z}} & 0 \\ 0 & e^{-\frac{1}{2}it\omega_{\Delta\bar{E}_z}} \end{pmatrix} \quad (5.9)$$

This is nothing more than the unitary evolution of states separated by the energy $\Delta\bar{E}_z = \hbar\omega_{\Delta\bar{E}_z}$.

A second rf pulse is then applied to complete the interferometer procedure:

$$\langle 1 | R_{\frac{\pi}{2}}(t + \tau, \tau) U(\tau) R_{\frac{\pi}{2}}(0, \tau) | 0 \rangle = \sin^2 \left[\frac{1}{2} t (\omega_{\text{rf}} - \omega_{\Delta\bar{E}_z}) \right]. \quad (5.10)$$

The only unknown parameter of the interferometer output is $\omega_{\Delta\bar{E}_z}$. Our measurement scheme is to trace the interference curve by scanning the rf frequency. From this scan we can measure the phase $\phi_{\Delta\bar{E}_z} = \omega_{\Delta\bar{E}_z} t$. We then repeat, say after reversing the field to extract a relative phase shift. This can then be used to measure any asymmetry in the electric field.

To simplify the calculation, we have made the approximation that the rf transition takes place on resonance, which obviously conflicts with our approach of scanning the rf frequency. Fortunately, the period of the Ramsey interference oscillation is very fast compared with the width of the rf transition. The result of this is that the final lineshape should in fact reflect both the form of the single loop rf lineshape, and the rapidly oscillating fringes of the Ramsey interference, defined in equation 5.10. Figure 5.13 shows such a Ramsey interference scan, with the characteristic rf lineshape envelope and Ramsey fringes clearly visible. We also need to make sure that the fringes shift by a small frequency upon reversal, compared with the fringe spacing (and of course, remain close to resonance), otherwise we could find we have asymmetric reversals which shift the




Figure 5.13: Ramsey interferometer fringes. This clearly shows the rapid oscillation of Ramsey fringes, under the envelope of the rf transition.




Figure 5.14: Ramsey fringes recorded at plate voltages of ± 4 kV, in each polarity state. Analysis of this data gives a difference in field of 31 ± 6 mVcm $^{-1}$. The curves are fits to $\alpha + \beta \sin(\omega V + \phi)$, an approximate model for the Ramsey interferometer output.

interferometer by multiple fringes. We know from the field mapping results that we can map fields to 0.6 Vcm $^{-1}$, whilst a Ramsey interferometer phase shift of π (0.5 kHz) at 20 kVcm $^{-1}$ is equivalent to 1.6 Vcm $^{-1}$, so combining both approaches we can reliably interpret the Ramsey interferometer fringes.

5.2.2 Ramsey interferometer results

Two Ramsey interferometer experiments were performed, one at each field polarity. Each time the rf frequency was scanned 3 kHz across the centre of the rf resonance. The resulting scans, along with fits, are shown in figure 5.14.

As can be seen the reversal is qualitatively very good, with the Ramsey fringes apparently in phase. By fitting the data to the interferometer lineshape, we see that the field magnitudes are slightly asymmetric, with a difference of 31 ± 6 mVcm $^{-1}$. In comparison with the Stark mapping technique presented in section 5.1.2, the precision with which we measure the average electric field has increased two orders of magnitude from ± 0.6 V to ± 6 mV. This technique offers improved sensitivity, though it is best suited to measuring relative fields, such as the quality of field reversals.

Chapter 6

EDM Results

This thesis does not report a new limit on the electron EDM. A robust measurement will necessitate a longer data set along with a complete analysis of systematic effects, null tests, and perhaps a control experiment using an insensitive molecule such as CaF. The result presented here are the analysis of a data set taken between March and May 2007. It should be viewed as a guide to current sensitivity, and an illustration of the analysis procedure.

6.1 EDM data analysis

The first step of the EDM analysis is to select blocks¹ to be included. The initial approach is to include every block, unless there is a *clear* reason not to. There are several motivations to exclude data. Firstly broken blocks where the laser unlocks are rejected. This is the most common reason for ending a cluster². Such events are easily identified by inspection of the raw data.

Another reason that blocks are not included is poor control of the magnetic field bias. As described in section 3.6.2, the magnetic field must be carefully set, such that the interference

¹Recall a *block* is a series of shots acquired in conjunction with appropriately switched fields to yield a single EDM measurement.

²Owing to the symmetry of the switching patterns, there is no value in part-completed blocks.




Figure 6.1: The block data consists of a series of integrated TOF profiles, which suitably summed yield the previously discussed analysis channels. *a*: A normal block with little drift. The points fall loosely into two levels due to the ΔB step. *b*: Data acquired as the laser unlocks. It is clear that after 1000 shots, the laser jumps off resonance. Such data is not included in the EDM analysis




Figure 6.2: Histogram of EDM data, before magnetic field veto, after damaged blocks are removed. Through comparison with a normal distribution of mean 0 and variance 1 (solid blue curve), it is clear that the normalised EDM data is not normally distributed.

fringes are centred upon the *DB* magnetic field step. Failure to centre the fringes leads to stepping about a part of the interference curve which may deviate strongly from the linear approximation used in the block analysis.

After rejecting damaged data 824 from a total of 1240 blocks remain. This fraction of rejected data is somewhat higher than we might ordinarily expect. During this particular data run, we were developing the field mapping techniques, which compromised the efficiency of running the EDM experiment, and fewer blocks were recorded each day.

Each 4096 shot block takes approximately 3 minutes to acquire, allowing for data analysis and plate charging and switching delays (see section 2.5.2). Sensitivities can be reported in $e.cm$ per $\sqrt{\text{block}}$, per $\sqrt{\text{day}}$ and per $\sqrt{\text{dataset}}$. We assume that in a perfect day of acquisition we acquire $\frac{24 \times 60}{3} = 480$ blocks.

After rejecting the blocks that were known to be damaged the EDM data must undergo a magnetic field veto: blocks acquired when the magnetic field was noisy are rejected, as described in section 3.7.3.

It is interesting to plot a histogram of normalised EDMs³ before the veto, shown in figure 6.2. This clearly shows that there is more weight in the wings of the EDM histogram than the normal distribution, which suggests the data is not normally distributed. For this reason the EDM confidence intervals are calculated using the bootstrap, as discussed in 3.7.3, rather than assuming gaussian statistics to derive a variance.

6.1.1 Bootstrap of EDM data before veto

In order to generate a confidence interval which better describes uncertainties, given the distribution of EDM data, the bootstrap technique is used (see section 3.7.3, as well as Efron [46]). The first step is to use the bootstrap to generate a cumulative density function. After the resampling 40000 replicate data sets from the original set of blocks, an approximation of the distribution of the measured EDM is generated.

³By normalising the histogram I mean that for each block, I divide the EDM by the EDM standard error of that block.

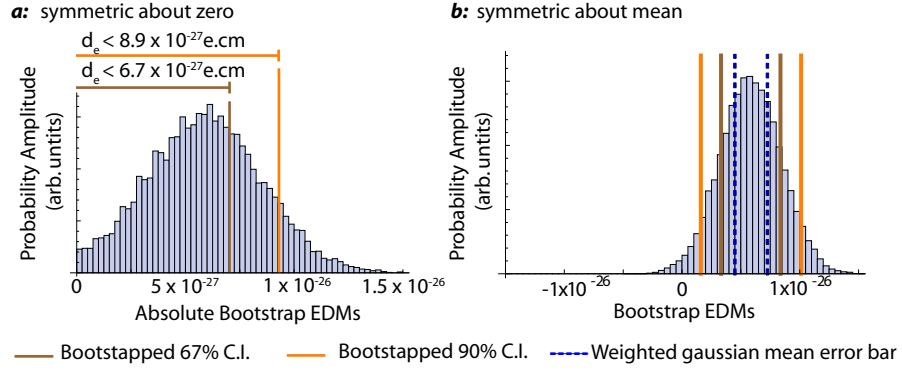


Figure 6.3: Confidence intervals for the measured EDM, before the magnetic field veto. *a*: The bootstrapped distribution about zero. *b*: The bootstrapped distribution about the mean. As no claim is being made about the central value of the EDM, the interval symmetric about the mean is the most appropriate to consider.

68%	90%
$3.3 \times 10^{-27} \leq d_e \leq 8.4 \times 10^{-27}$	$1.6 \times 10^{-27} \leq d_e \leq 10.1 \times 10^{-27}$
$ d_e \leq 6.9 \times 10^{-27} \text{ e.cm}$	$ d_e \leq 9.1 \times 10^{-27} \text{ e.cm}$

Table 6.1: Symmetric and absolute Bootstrapped EDM confidence intervals, before magnetic field noise veto.

For completeness, I consider two types of confidence interval (CI). The first is constructed symmetrically about the mean of the bootstrapped distribution, describing the interval *lower bound* $< d_e < \textit{upper bound}$. For completeness, a second interval is also calculated. This is constructed as the smallest interval symmetric about zero which contains some fraction of the distribution. It provides the absolute bound $|d_e| < \textit{Absolute bound}$. This interval is most applicable to data which is thought to be consistent with zero, as it is similar to the above interval in the case that the mean is zero. To do this a cumulative distribution function (CDF) is generated based upon the absolute value of each bootstrapped EDM replicate. The calculated confidence intervals for both 68% and 90% intervals are shown in figure 6.3 and table 6.1.

The bootstrapped distribution lies some way from zero, and as a result the confidence interval symmetric about zero is somewhat wider than the mean centred confidence interval. We also note, as we should expect from the histogram shown in figure 6.2, the simple variance estimator suggests a more sensitive measurement than the equivalent 68% asymmetric confidence interval.

The next step is to use the magnetometer data to reject data taken when the magnetic field noise was very high.

6.1.2 Magnetic field veto

In section 3.7.3 we saw that the magnetic field noise can limit our sensitivity. This happens when, expressed as an EDM, the Zeeman shift is comparable to the size of the block's EDM error bar. The veto procedure compares these two quantities and excludes blocks where the magnetic field




Figure 6.4: Investigation of magnetic field noise veto level. *figure a*: As the severity of the veto is increased fewer blocks are included. *figure b*: Comparison of the bootstrapped EDM error bar to the modified standard deviation across the retained data set.

noise was high. Decreasing the veto level leads to more aggressive data rejection, and in turn the inclusion of fewer blocks.

The motivation for doing this is that in the case where the EDM data is not normally distributed, but has substantial weight in the wings of the distribution, our result can be unexpectedly influenced by outliers taken when the magnetic field noise was dominant.

An analysis of the EDM data with progressively relaxed vetoing is shown in figure 6.4. As the veto became more severe, the number of included blocks included was reduced (figure 6.7.a). In figure 6.7.b I investigate the EDM error bar. Shown here are two estimators - the modified standard deviation, is simply the error bar derived using gaussian statistics, as previously described in section 3.6.3.

It is clear that as the veto is relaxed, increasingly noisy blocks are included in the analysis. In doing so, the bootstrapped confidence interval width can be seen to decrease, then at approximately 5, plateau. Also evident, as discussed earlier, is that the bootstrap analysis better reflects the additional weight in the wings of the distribution - leading to a wider confidence interval. Given this, I re-analyse the EDM data with a magnetic field veto of 5.

The effect of the veto process is shown in figure 6.5. The veto level of 5 rejects blocks with signal - magnetic field noise ratios up to 14 times greater than the veto level. Condylis noted in his research [14] that the afternoon and early evening were particularly poor times for taking data, as the magnetic field noise was high. From the analysis of the data taken here it is not possible to carry out such an analysis, as data was solely taken during normal laboratory operating hours. Normalising against the frequency of blocks taken each hour, the rate of block rejection appears reasonably uniform throughout the day.

Analysis of EDM data with magnetic field veto

If we use the magnetic field veto at 5 we lose nearly one fifth of the data, from 824 to 677 blocks. The resulting EDM analysis is $d_e = 2.8 \pm 1.65 \times 10^{-27}$ e.cm. Again, the histogram of each blocks'




Figure 6.5: Plot of veto figure of merit against time. As may be seen, the veto rejects blocks with magnetic noise/EDM signal ratios very much greater than 5, the veto level. The red points (above the dashed line) represent rejected blocks, the blue points are included the EDM analysis. From this figure it is also apparent that most data was taken during the afternoon/early evening, a time Condylis [14] concluded to be the noisiest time of the day. Unfortunately, no data was taken through the night to compare statistics.




Figure 6.6: EDM data analysis. Histogram of vetoed block EDMs, along with normal distribution. As with the retained data, the data is not normally distributed.

evaluated EDM⁴ is compared to a normal distribution in figure 6.6.

Inspection of the histogram suggest the data is not normally distributed. Computing bootstrapped confidence intervals as before demonstrates slightly wider bounds (table 6.2).

68%	90%
$1.2 \times 10^{-27} \leq d_e \leq 5.5 \times 10^{-27}$	$-1.6 \times 10^{-27} \leq d_e \leq 7.4 \times 10^{-27}$
$ d_e \leq 4.0 \times 10^{-27} \text{ e.cm}$	$ d_e \leq 6.3 \times 10^{-27} \text{ e.cm}$

Table 6.2: Bootstrapped EDM confidence intervals, after veto.

The confidence intervals are shown together with the bootstrap replicate distribution in figure 6.7.

6.2 Discussion of results

The most important number in this thesis is the sensitivity at which we can operate. We see that after the bootstrap, the remaining blocks give us a 68% confidence interval half width of $2.15 \times 10^{-27} \text{ e.cm}$. Although this data was taken over three months, it was not a determined data

⁴Normalised as μ/σ .




Figure 6.7: Confidence intervals for the measured EDM, after the magnetic field veto. These were calculated using a 40000 replication bootstrap.

run, as we were developing the mapping scheme during this time: 677 blocks in 3 months is not the limit of our rate of data acquisition.

A useful measure of our experiment is our sensitivity per $\sqrt{\text{block}}$ and per $\sqrt{\text{day}}$. Our total data set is 824 blocks long, 677 of which were used. We might also consider the situation where we improve our efficiency, such that no data need be rejected, perhaps through better noise rejection or magnetic shielding⁵.

	Perfect (no rejected data)	Typical
sensitivity/ $\sqrt{\text{block}}$	$5.6 \times 10^{-26} \text{ e.cm}$	$6.2 \times 10^{-26} \text{ e.cm}$
sensitivity/ $\sqrt{\text{day}}$	$2.55 \times 10^{-27} \text{ e.cm}$	$2.84 \times 10^{-27} \text{ e.cm}$

Table 6.3: Sensitivities of the EDM experiment, after magnetic field veto.

An interesting calculation is how much more integration would we need to have a sensitivity equal to that of Commins ($7.4 \times 10^{-28} \text{ e.cm}$). Taking our typical running sensitivity, we would need 11 times as much data, approximately 7500 blocks (after veto). Given that further data must be rejected when the laser unlocks, this figure will likely rise to 10000. Running efficiently, 24 hours/day this suggests we could integrate to Commins' level of sensitivity in around three weeks⁶.

Given this, why did we not carry on? Ultimately, with the instrumentation and systematic rejection at the time of taking this data it was felt that more work was required to support our measurement, hence the development of the field mapping techniques documented in chapter 5. Given accurate maps of the fields in the machine, improved models of the interferometer can now be constructed, which in future will be used to exhaustively search for all possible systematics. Of course, completion of all null and control tests must also be taken into account. Though null tests can generally be run unattended⁷, the CaF control experiment will take a similar amount of

⁵In fact, in a 24 hour/day data acquisition program we would expect the magnetic field noise to be far less at night.

⁶This is actually a little pessimistic. Gains in sensitivity will be made running at night, due to reduced magnetic field noise. Additionally, running for long periods of time allows better source optimisation.

⁷Null tests without a molecular beam may also be run very much more rapidly as the PMT gates are set very much wider than usual, greatly improving speed at which null tests may be performed.




Figure 6.8: Comparison of veto and non-vetoed bootstrap distributions. To generate the PDF, means and variances were fitted to the CDF (a). Given these, the overlap of the PDFs was calculated to be 0.54 (b)

time to complete. Finally, we must also undertake experiments at a range of voltages, to ensure any $E.B$ channel analysis scales appropriately with the polarisation η . This complete set of tests would require modification to instrumentation and our running procedure before the experiment can be considered sufficiently robust. These are detailed in the final chapter of this thesis.

Effect of the veto on the EDM measurement

The magnetic field veto was introduced to discard the noisiest blocks - blocks where the magnetic field variation was so strong that the interferometer output is likely to be strongly influenced. In the single data set shown here it has - perhaps by chance - moved the central value of the EDM analysis towards zero. In this section I will examine whether this trend is statistically significant.

In the first instance, I'll compare the raw EDM data to the retained data (veto level of 5). This process rejected 18% of the data. In doing so the central value moved from 5.6×10^{-27} e.cm to 2.8×10^{-27} e.cm. The concern is that perhaps there is systematic shift in the EDM associated with the most noisy blocks? To evaluate this I take each data set and resample a set of 10000 bootstrapped replicates. The resampled distribution approaches a normal distribution so fitting a mean and variance to the CDF (6.8.a) reasonably describes the distribution. A comparison of the fitted normal probability distribution functions is shown in figure 6.8.b.

Though the veto does, with this data set, move the central value, it is not statistically significant given the uncertainties considered. The post-veto bootstrap distribution has significant overlap with the bootstrapped full data set.

Though this approach seems appropriate, it is somewhat misleading: the two distributions are not independent. As the veto rejects less than fifth of the data, more than 80% of the data is common to each distribution. It should come as little surprise that the PDFs appear similar. Instead in figure 6.9 then we compare the bootstrapped distributions of the included and rejected data after applying the veto.

Again, in this analysis, a shift in the mean is apparent, though the overlap of these PDFs is




Figure 6.9: Comparison of rejected and retained bootstrap distributions. As before, bootstrap CDFs for each set of blocks were generated and fitted to a normal distribution CDF. Taking the fitted mean and variance, PDFs of a normal distribution were produced for each data set.

much smaller (14%).

A useful measure is to compare the difference in means ($\mu_{\text{rejected}} - \mu_{\text{retained}}$) of these distributions with the sum of the variances ($\sqrt{\sigma_{\text{rejected}}^2 + \sigma_{\text{retained}}^2}$). These quantities compare as $12 \pm 6 \times 10^{-27}$ e.cm. The error on the difference is of order the difference in means.

In all it is not altogether clear that the veto systematically shifts the central value of the EDM analysis. It will be necessary to monitor this effect in a larger data set. If the noisy data truly changes the central value of the EDM analysis, then the difference ($\mu_{\text{rejected}} - \mu_{\text{retained}}$) will likely be similar, however with more blocks in a larger data set, the error bars should be much smaller, and allow a more informed conclusion to be drawn.

The effect of the veto on the EDM analysis can be interpreted in two ways. Either there is a persistent trend that the data acquired when the magnetic field was noisy dominates the central value, or we have insufficient samples to draw a useful conclusion, and that more data, noisy or otherwise, would converge to zero.

If at times there is a large change in magnetic field correlated with the state of the electric field, then we have a candidate for an EDM. However, we can detect and reject this. Outside the magnetic shields, the flux gate magnetometer is sufficiently sensitive to measure magnetic fields - expressed as EDMs - much smaller than 10^{-28} e.cm. The magnetic field veto can exclude blocks dominated by this effect.

One situation remains that must be considered, that there is a magnetic field originating from inside the inner shield. In this case it is unlikely that the magnetometer could reject a systematic effect. By far the most likely cause of such an effect is a current associated with the electric field plates. In this case we have several defences. Firstly current monitors could be used to sample charging / leakage currents throughout the entire data acquisition process (see §7.2.1). Secondly we might also monitor how any measured EDM scales with voltage. A true electron EDM must scale as the polarisation of the molecule - figure 1.5. If the EDM were to be associated with a

charging current, then we might expect it to scale linearly with voltage⁸. If associated with electric breakdown we would expect a very non-linear scaling, most likely with an EDM only obvious at the high electric field.

⁸This reasoning is based on magnetic fields associated with leakage currents. As capacitive charging currents scale linearly with voltage, so to should any generate fields that the interferometer may sample.

Chapter 7

Conclusion

7.1 Main thesis results

The EDM experiment has been evolving over a number of years, previously documented by Redgrave [44], Hudson [16] and Condylis [14]. Each generation of experiment builds upon the next. Hudson's thesis describes the early implementation, broadly similar to that presented here. Condylis implemented an improved source, moving from an effusive oven to the pulsed molecular beam used in this thesis. He also implemented Raman transitions as a means of splitting/recombining the interferometer. The work presented in this thesis is the culmination of a series upgrades that have been introduced to better constrain systematic effects. Experimentally, the most obvious change is in the how we perform rf transitions. Whereas Hudson used long rf pulses, and Condylis used Raman transitions, in this thesis short rf pulse transitions have been developed. This has lead to more localised transitions, which in turn sample less inhomogeneous fields. Other important modifications were the move from triple to single region electric fields, and from rf loops the rf transmission line. We have also developed new methods of mapping of electromagnetic fields within the interferometer.

The overall sensitivity of the experiment has not improved. Though improved pumping and a faster ablation laser should have been straightforward upgrades. In fact, even with improved pumping (§2.0.1) our signal at 50 Hz was very similar in quality to that at 25 Hz. Target life was reduced, so there was no overall gain. Condylis achieved a sensitivity of $1.51 \times 10^{-27} \text{ e.cm}/\sqrt{\text{day}}$. Before bootstrapping, but still using the magnetic field veto, this thesis reports a sensitivity of $2.55 \times 10^{-27} \text{ e.cm}/\sqrt{\text{day}}$ (meaning we need roughly 2.8 times as much data to reach similar sensitivities). These results are only loosely comparable as one can achieve arbitrarily high sensitivities by vetoing many more noisy blocks¹. Furthermore, it should be considered that Condylis took some data at night, when magnetic field noise is known to be reduced.

¹The correct measure must be the sensitivity of blocks remaining after the veto, scaled by the total number of blocks acquired.

7.2 Future progress

Future experimental progress can be organised into two groups - upgrades required before acquiring a robust EDM data set, and long term projects to improve sensitivity. These are illustrated in figure 7.1 below. Immediate upgrades to facilitate a competitive measurement are likely to be completed by early 2008. The longer term development of the experiment is discussed below.

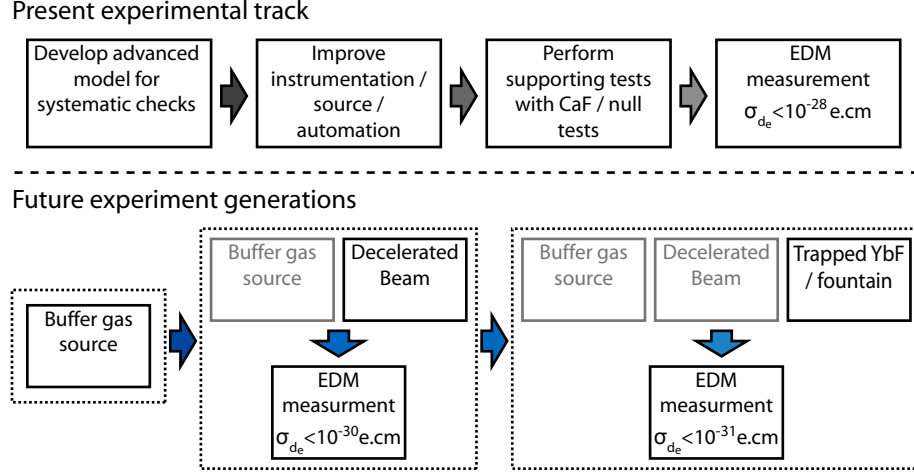


Figure 7.1: Future development of the EDM experiment. In the short term, with more development of the systematic models, the current experiment will probe to $\sigma_{de} < 10^{-28} \text{ e.cm}$. Beyond this several generations of development build towards a trapped YbF EDM experiment, with very much longer coherence times.

7.2.1 Short term development

In order to support the EDM measurement a series of instrumentation upgrades have been planned. These focus on controlling electric, magnetic and rf fields. Primarily these move the experiment towards autonomous operation, but also allow for a more precise understanding of the experimental environment.

Systematic effects

An important task to be addressed before measuring an EDM is to improve our understanding of systematic effects. I have previously discussed a new systematic effect, associated with the old three region field design which has been suppressed. We have no evidence to suggest that we have any remaining systematic effects, though obviously there remains that effects may be noticed as we improve out sensitivity.

Despite admitting that an exhaustive search for systematic effects is impossible, we are still in a strong position to make a robust measurement. By combining EDM measurements of different molecules at different fields we can ensure that our measurement is free from systematic effects, albeit at the expense of taking several data sets.

Eventually a more detailed model could be used to improve our understanding of possible non-obvious systematic effects. These include the effects of magnetic fields present in the rf transition region, a higher order expansion of the interferometer lineshape and the effects of taking data with poorly centred interference fringes.

Electric field control

A more advanced scheme for electric field control has been designed. Recall the current scheme uses a 0 V - 5 V voltage source to control a 50 kV Bertan HV power supply. A digital line is used to ground this reference to turn off the field as required². Though the digital control of field state is adequate, manual voltage control complicates experimental running. Firstly, the initial voltage must be set as accurately as possible by hand. This is time consuming, and requires very disciplined operation. Furthermore, over the timescale of a cluster, the potentials have been seen to drift several volts, especially when power supplies are warming up. A computer control system could reset the monitor and ‘lock’³ throughout the cluster. Our improved scheme will use an optically isolated National Instruments USB-6215 multi-IO, where an analogue output is to be used to provide the voltage reference for the HV supply. We plan to monitor the power supply output using the existing divider and an analogue input.

This scheme brings several advantages. Most obviously we can now set and monitor the voltage using the computer. Once the calibration of the HV divider has been established, the applied potential may be continuously monitored.

Another interesting application of this control is to step the voltage a small amount⁴ using a spare switching pattern. By looking for an EDM correlated to this waveform we can probe systematic effects due to imprecise plate potentials. For example, considering the electric field systematic discussed in section 4.2.1 (effects associated with multiple electric field regions), we believe after calculation that the move to single region plates suppresses the effect. However, by introducing a small switched asymmetry into the applied potentials we could monitor this effect in our final data set.

Finally, we can optimise our electric field reversals. Currently to reverse the electric field polarity, we set the power supply control voltage to zero, and allow the plates to discharge through a resistor into the power supply. When measuring the EDM we also have to allow time for the charging currents to decay. We also try to minimise the discharge current to avoid magnetisation of any nearby ferromagnetic material. Typically the plate charge decays exponentially over 10-15 seconds. This scheme is neither fast nor free from introducing large discharge currents. However, for most of the discharge time the leakage current is small. Ideally we would like to optimise the

²The electric field hardware is described in figure 2.2.

³Though we could servo the electric field voltage, this seems like an uncomfortable approach, that may admit very complicated systematic effects. A simpler, solution is to reset the electric field every block.

⁴Not sufficient to move the rf transitions appreciably off resonance.

charging/dischage such that the leakage current is constant, and below a predetermined ‘safe’ level. By ramping the applied voltage appropriately we can achieve a constant charging current, with much faster reversals, and remove the need for series charging resistors. Reducing the switching time is beneficial, as we can increase the E switching pattern frequency, moving the E-switching channel away from the $1/f$ part of the noise power density spectrum, without worsening the duty cycle.

Computer controlled magnetic field calibration

As the magnetic field in the lab slowly drifts the phase of the interferometer fringes change, which eventually leads to the interferometer operating away from the ‘linear’ part of the curve, possibly rendering data susceptible to systematic effects. This phase is observed in the B-channel. At present, the data acquisition must be closely monitored, such that this is manually corrected as necessary - usually several times each day. In order to allow for autonomous data acquisition, we need the control software to correct for this automatically. The planned procedure trims the B-field offset, based upon the analysis of the EDM data. This is a relatively minor upgrade, using a spare analogue output to control the magnetic field bias current.

Frequency doubled fibre laser

The main pump/probe laser beams originate from the Spectra Physics 380D dye laser. Though capable of providing the light we need, this laser is very labour intensive to run. The first drawback is the combined time to allow the laser to stabilise (1 hour per day), keep the laser locked (regular tuning) and periodically replace the dye (1 day / 2 weeks). Perhaps more significant is that this laser prevents us from running unattended.

The new laser requirements are stability and linewidth, as discussed in section 2.1.2. The main reason for an upgrade is to reduce operating complexity. An ideal laser would be easier to lock, more robust, and require less servicing. Unfortunately, at 553 nm we are very limited in our choice of laser. Recently though the commercialisation of appropriate fibre lasers has allowed us to develop a new laser system; a frequency doubled fibre laser.

The IR source is a 1106 nm ytterbium doped fibre laser. The frequency of the laser is set at production, by etching an appropriate grating into the end of the fibre. Thermal tuning allows small adjustments of 0.7 nm in addition to a piezo in the seed laser which gives very fine control. The laser produces up to 1 W of IR.

The IR is then frequency doubled in a lithium triborate (LBO) crystal. We need a reasonable amount of doubled light (10 – 50 mW), so a single pass doubling scheme will not suffice⁵. To increase conversion efficiency, the LBO is placed in a stabilised ring cavity. Even with moderately low finesse, this should lead to the generation of 50 – 150 mW green.

⁵A single pass scheme using periodically poled Lithium Niobate was attempted, but proved inefficient.

The real advantage of this scheme is the long term stability. To find the correct transition, the fibre laser frequency must be tuned over the main probe transition. This can be carried out using the computer, just as we scan any parameter space. Fitting the lineshape to an appropriate model, we can then set the piezo to the correct voltage. In the simplest case this could be carried out periodically during data acquisition to keep the laser locked. The specification of the fibre laser suggests a sufficiently low drift rate that the laser should be stable over the timescale of a cluster. A more advanced scheme is to dither the laser frequency either side of the line centre as we are taking EDM data, controlled by another switching pattern. This channel can then be analysed to calculate how far the laser is from resonance, just as we can set the magnetic field bias by using the B -channel. We could also probe systematics: if the laser frequency dither is the ν -channel, we can analyse $\nu.E.B$, to see if the EDM varies with laser frequency.

Automated data acquisition

Operating the EDM hardware requires human attention, so to run 24 hours per day is a serious undertaking, unsustainable over long periods of time. However, many of the tasks involved in operating the experiment are straightforward, and possible to automate. Certainly the most difficult task to overcome is locking of the dye laser. However, with the implementation of the fibre laser source there remain few technical challenges:

- E field monitoring/control
- Leakage current monitoring
- Target rotation
- Magnetic field bias correction
- Automatic rf optimisation
- Beam optimisation

Given the improvements suggested above, these upgrades are straightforward. Automatic rf optimisation requires minor upgrades to the EDM data acquisition software, to record and fit a scan of rf power and rf frequency.

Leakage current monitoring is a relatively straightforward, but valuable upgrade. New leakage monitors have been designed with 1 nA sensitivity. These encode current information in a frequency modulated square wave, which is transmitted optically over fibre to a National Instruments counter. The computer can then continuously monitor leakage and charging currents throughout data acquisition.

Target rotation requires the minimal upgrade of a stepper motor to increment the target position - moving a ‘fresh’ section of Ytterbium in front of the ablation laser.

Beam optimisation is a little more complicated. The software controllable parameters are the timings (relative to the Q-switch of the Nd-YAG ablation laser) of the Nd-YAG flashlamps, gas valve opening time, and duration of the gas valve opening. The simplest way to optimise the beam is to maximise each of these independently. These parameterisations change as the target degrades, and vary somewhat from day to day. The simple solution might be to leave them fixed, and to only adjust the source position when needed. However, given that with automation much more data can be acquired, we may be able to fully explore the source parameter space and develop a more sophisticated heuristic.

Additionally under an automated scheme we expect data taken at night to be more sensitive. This is due to reduced magnetic field noise outside of the physics department’s normal operating hours. This was first observed by Condylis ([14], figure 6.10).

Automated data acquisition would allow for near ‘perfect’ data acquisition. In such a regime, we could take 1000 blocks over a weekend. It is a reasonable assumption that automated running of the experiment, and making measurements at night we will reach and likely improve on Condylis’ sensitivity of $1.51 \times 10^{-26} \text{ e.cm}/\sqrt{\text{day}}$. Doing so will mean we need less than 5 days data.

7.2.2 Long term development

In the long term we are interested in improving our sensitivity by several orders of magnitude. Assuming gaussian statistics for a given sensitivity/ $\sqrt{\text{block}}$, 100 times more data is required for an order of magnitude increase in sensitivity. Clearly great gains cannot be made through extended data acquisition. Having said that, the convenience of a fully automated data acquisition scheme is likely to bring about a factor of 5 improvement in sensitivity⁶.

Using an interferometer scheme, the EDM is measured through the phase $\phi_{d_e} = d_e E_{\text{eff}} \eta t / \hbar$. Higher effective electric fields may exist in other molecules (see Tarbutt [50]), but our substantial investment thus far, combined with limited gains makes it more practical to continue using YbF. The polarisation of the molecule η is typically 0.7, so order of magnitude gains cannot be sought with higher electric fields. To improve our sensitivity we should first look to increasing the time spent in the interferometer.

Decelerated YbF

Experiments to decelerate polar molecules are being carried out by Tarbutt ([51]). The general concept is to use the Stark shift of the molecules together with a dynamically switched electric field to provide a deceleration of the molecular pulse. The method used was an alternating gradient, strong field seeking decelerator. Strong field seeking molecules are attracted to a high electric field. The concept was to arrange a series of electrode pairs along the beam line. Each pair was switched

⁶A determined data run longer than a week is not feasible with manual data acquisition. However, using automation could bring about a factor of 25 increase in data volume.

on just as the molecular pulse passed, such that the pulse experienced a series of small decelerating forces. At present, a 12 stage decelerator has been shown to remove 12% of the molecules' (YbF) kinetic energy. Unfortunately, due to the quadratic scaling of kinetic energy this corresponds at present to a fairly small change in velocity - from 500 ms^{-1} to 482 ms^{-1} .

The choice of strong field seeking states causes some problems. As Maxwell's equations prohibit a field maxima in free space, it is impossible to contain the molecules in the transverse plane using static fields. Dynamically alternating the gradient of the field, using offset electrodes in principal allows for refocussing, but electric field edge effects and the spatial extent of the molecular pulse lead to inefficiencies. These ultimately limit the number of stages of the decelerator, and hence the possible total deceleration.

An improved 'weak field' decelerator is currently being constructed, where in weak field seeking states the molecules are drawn to regions of low electric field. In this scheme there is a limit on the maximum electric field which may be applied, so each stage of the decelerator is less effective. However, it is possible to create a static field minimum in free space, hence containing the molecules along the beamline. Because fewer molecules are lost, there is no limit to the number of stages that may be incorporated and deceleration to rest should be possible.

Buffer gas source

The molecules created by the current supersonic expansion source have high velocity. This limits our present sensitivity due to the short 1 ms coherence time, and also makes deceleration challenging.

The buffer gas source [52] uses an aluminium fluoride - ytterbium target, encased in a chamber of cool (3 K), high pressure (1 Torr) helium buffer gas. As with the supersonic source, a Nd-YAG laser is used to ablate the surface of the target, leading to the formation of YbF. The YbF rapidly thermalises with the surrounding He, before effusing through a small aperture into a larger vacuum assembly. The characteristics of this source compare favourably with the supersonic source. The buffer gas beam produces pulses of velocity 50 ms^{-1} - 200 ms^{-1} , compared the supersonic pulse velocity of $\sim 500 \text{ ms}^{-1}$. This in itself represents a huge improvement of up to an order of magnitude velocity reduction. This may be more important in the context of deceleration, as slowing molecules requires a quadratic reduction in energy - deceleration from 500 ms^{-1} to 50 ms^{-1} accounts for 99% of the energy required to bring the molecules to rest.

A promising scheme to further improve sensitivity is to implement a fountain experiment, where coherence times may be up to 3 orders of magnitude longer than at present. This experiment cannot be implemented until YbF has been slowed significantly.

7.3 Outlook

The long term improvements can be incrementally introduced into a series of experiments. The buffer gas source alone could bring about a 10 fold reduction in beam velocity, such that we might gain an order of magnitude in sensitivity. Combining the decelerator and buffer gas source will further reduce the beam velocity. At this stage we will either make a further low velocity measurement, or implement a fountain type experiment.

Future generations of the experiment look likely to probe the electron EDM well beyond the present experimental limit. As stated in the introduction, Commins' measurement, the present experimental limit, already places constraints on values predicted by supersymmetry. A future measurement with a statistical uncertainty of $\sigma_{d_e} < 10^{-30} \text{ e.cm}$ would certainly either verify or confute theory.

Appendix A

Experiment parameters

A.1 Molecular beam optimisation

At the time of writing, the following parameters can be used as a guide to optimising the EDM experiment. Rather than specify a specific value for each parameter, I quote ranges to scan. The valve timings should be set to whichever value gives the greatest integrated TOF. The Nd-YAG should be optimised as discussed in figure 3.2.

Parameter	Control	Comment
	Typical values ^a	Lab book notation
Valve Pulse Length	Timing edge (PG)	Optimise with scan ^b
	120 μ s - 200 μ s	<i>VPL</i>
Valve to Q	Timing edge (PG)	Optimise with scan
	300 μ s - 350 μ s	<i>VTQ</i>
Flash to Q	Timing edge (PG)	Optimise with scan
	200 μ s - 250 μ s	<i>FTQ</i>
Valve voltage	Manual control	Monitor TOF ^c
	325 V - 375 V	n/a
Valve tension	Manual control	Optimise with FIG after moving source
	n/a	n/a
Valve alignment	Manual control	Optimise with FIG after moving source
	n/a	n/a
Yb surface quality	Check periodically ^d , rotate as required	Resurface using lathe or Scotchbrite
	n/a	n/a
Pump/probe beam alignment	Manual control	Adjust position to gain maximum signal/noise
	n/a	n/a
Pump/probe beam power	Manual control	Set to desired power using EDM expt. power meter ^e
	Pump 600 mV ₃	n/a
	Probe 900 mV ₃	
Nd-YAG focus/alignment	Manual control	Adjust as described in text
	n/a	n/a

^a Where appropriate, some of these items are just adjusted qualitatively to maximise the LIF/FIG signal

^b Depends strongly upon valve voltage. Avoid higher voltages combined with longer pulses.

^c When scanning the valve voltage by hand monitor the TOF to look for irregular pulse shape. Monitor chamber pressure and poppet sound.

^d Check the surface of the Yb whenever the source fails to optimise, especially after extensive use.

^e Pump/probe powers are measured using a Photodiode connected to a set of potential dividers. All voltages are measured on scale '3'.

Table A.1: Main EDM experiment parameters

Appendix B

Theoretical description of rf pulses

Central to understanding the interferometer scheme is the theory of how the YbF interacts with an applied oscillatory magnetic field. I start here by following Ramsey's[49] derivation the interaction of a two level system with an oscillating perturbation. The four level system can be treated in a similar manner, though the increased number of states makes the algebra complex.

At time $t = 0$, an oscillatory perturbation is applied which leads to transitions between to states p and q . This leads to the perturbation $V_{p,q}$

$$V_{p,q} = \int \psi_p^* V \psi_q d\tau = \hbar b e^{i\omega t}, V_{q,p} = \hbar b e^{-i\omega t}, \quad (\text{B.1})$$

where b is a measure of the strength of the interaction.

Given that the system is confined to the states p and q , with respective energies W_p and W_q , we can write the wavefunction which describes the system as the sum

$$\Psi(t) = C_p(t)\psi_p + C_q(t)\psi_q. \quad (\text{B.2})$$

The time - dependent Schrödinger is then used to calculate the evolution of these total system

$$i\hbar \frac{\partial \Psi}{\partial t} = \mathcal{H} = \mathcal{H}_0 + V \quad (\text{B.3})$$

A pair of coupled differential equations can now be generated by multiplying the Schrödinger equation ψ_p or ψ_q . By considering orthogonality we find

$$i\hbar \dot{C}_p(t) = W_p C_p(t) + \hbar b e^{i\omega t} C_q(t) \quad (\text{B.4})$$

$$i\hbar \dot{C}_q(t) = W_q C_q(t) + \hbar b e^{-i\omega t} C_p(t) \quad (\text{B.5})$$

These equations can then be solved to yield expressions for the amplitude coefficients $C_p(t)$,

$C_q(t)$. Taking the case where the molecule is initially in state p , such that $C_p(0) = 1$, Ramsey calculates:

$$C_p(t) = \left(i \cos \Theta \sin \frac{1}{2}at + \cos \frac{1}{2}at \right) \exp \left\{ i \left[\frac{1}{2}\omega - (W_p + W_q)/2\hbar \right] t \right\} \quad (\text{B.6})$$

$$C_q(t) = i \sin \Theta \sin \frac{1}{2}at \exp \left\{ i \left[-\frac{1}{2}\omega - (W_p + W_q)/2\hbar \right] t \right\}, \quad (\text{B.7})$$

where $\cos \Theta = (\omega_0 - \omega)/a$, $\sin \theta = -2b/a$, $a = [(\omega_0 - \omega)^2 + (2b)^2]^{\frac{1}{2}}$ and $\omega_0 = (W_q - W_p)/\hbar$. Given that we start in state p , with $C_p(0) = 1$, the probability of the transition $P_{p,q}$, from state p to state q after time t , must be

$$P_{p,q} = |C_q(t)|^2 = \sin^2 \Theta \sin^2 \frac{1}{2}at \quad (\text{B.8})$$

$$= \frac{(2b)^2}{(\omega_0 - \omega)^2 + (2b)^2} \sin^2 \left\{ \frac{1}{2} [(\omega_0 - \omega)^2 + (2b)^2]^{\frac{1}{2}} t \right\}. \quad (\text{B.9})$$

This important result is the single pulse rf lineshape for a two level system. It is somewhat limited however, as to develop a double pulse ‘Ramsey’ interferometer, we need to understand the effect of an rf pulse upon an arbitrary state. Using matrix notation, the state of the system can be described by the vector $S(t) = (C_p(t), C_q(t))^T$. The general solution to the coupled differential equation B.6 is then

$$R = \begin{pmatrix} e^{i\tau(\frac{\omega}{2} - \frac{W_p+W_q}{2\hbar})} (\cos [\frac{a\tau}{2}] + i \cos[\Theta] \sin [\frac{a\tau}{2}]) & i e^{it\omega + i\tau(\frac{\omega}{2} - \frac{W_p+W_q}{2\hbar})} \sin [\frac{a\tau}{2}] \sin[\Theta] \\ i e^{-it\omega + i\tau(-\frac{\omega}{2} - \frac{W_p+W_q}{2\hbar})} \sin [\frac{a\tau}{2}] \sin[\Theta] & e^{i\tau(-\frac{\omega}{2} - \frac{W_p+W_q}{2\hbar})} (\cos [\frac{a\tau}{2}] - i \cos[\Theta] \sin [\frac{a\tau}{2}]) \end{pmatrix} \quad (\text{B.10})$$

This allows the evolution of the rf superposition to be calculated after the application of an rf pulse of length τ , applied at time t , $S(t+\tau) = R.S(t)$. Before discussing the four level system, consider a two level Ramsey interferometer, as used for measuring the average electric field throughout the interaction region (see section 5.2). As discussed in chapter 1, this is in many respects similar to our 4 level interferometer.

The first step of the interferometer scheme is to prepare the system in state p , or $S(0) = (1, 0)^T$. Assume this has been achieved, and there is no residual population in state q .

The first rf transition is achieved by tuning the applied rf field to $\omega = \omega_0$, and adjusting the interaction strength b , and the rf pulse length (τ), such the molecules are driven $\pi/2$ around the Bloch sphere, into the superposition $S(0) = \frac{1}{\sqrt{2}}(1, 1)^T$. For this, the rf interaction strength must be set to $b = \pi/4\tau$. Equation B.10 then reduces to

$$R_{\frac{\pi}{2}}(t, \tau) = \frac{1}{\sqrt{2}} \cdot \begin{pmatrix} e^{\frac{1}{2}i\tau\omega_0} & ie^{\frac{1}{2}i(2t+\tau)\omega_0} \\ ie^{-\frac{1}{2}i(2t+\tau)\omega_0} & e^{-\frac{1}{2}i\tau\omega_0} \end{pmatrix}, \quad (\text{B.11})$$

where the time dependance is introduced to allow for the correct interaction of the rf synthesiser phase with the phase of the state.

This derivation of a $\pi/2$ pulse assumes can be applied to any state S , so is used again for recombining after the free evolution stage. The free evolution of the system is simply the evolution in the case where $b = 0$:

$$U(t) = \begin{pmatrix} e^{\frac{1}{2}it\omega_0} & 0 \\ 0 & e^{-\frac{1}{2}it\omega_0} \end{pmatrix} \quad (\text{B.12})$$

The key to the EDM experiment is of course that we are trying to measure ω_0 with great sensitivity. With the two level system, it is the Stark shift, not the Zeeman or EDM terms that is exposed, but the principle is the same.

The second rf pulse is identical to the first, except of course that time has passed, allowing the relative phase of states p and q to evolve.

So, the two level Ramsey interferometer can be succinctly described by a series of state operations upon the initial state vector $S(0)$:

$$S_{\text{Final}} = R_{\frac{\pi}{2}}(t + \tau, \tau).U(\tau).R_{\frac{\pi}{2}}(0, \tau).S(0) \quad (\text{B.13})$$

The final stage is to measure the population of either state p or q , to expose the difference in phase accrued during free evolution.

Moving to a four level system, I take the following from Hudson's thesis [16], which is an extension of Ramsey's analysis of a two level system to a four level system.

Consider the hyperfine levels of the YbF ground state. These are described in the z - basis by state vector representing the $|F, m_F\rangle$ amplitudes $|0, 0\rangle, |1, -1\rangle, |1, 0\rangle, |1, 1\rangle$. A more suitable basis for calculating field interactions is used, which we call the cartesian basis. This consists of the states $|0, 0\rangle, |x\rangle, |1, 0\rangle, |y\rangle$, where $|x\rangle = \frac{1}{\sqrt{2}}(|1, 1\rangle + |1, -1\rangle)$ and $|y\rangle = \frac{1}{\sqrt{2}}(|1, 1\rangle - |1, -1\rangle)$. To transform between these basis, following transform is used:

$$\mathbb{CB} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ 0 & 0 & 1 & 0 \\ 0 & \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \end{pmatrix} \quad (\text{B.14})$$

Similarly to the two level system, a set of four coupled differential equations can be derived by considering interactions between all four states of the $F=0,1$ hyperfine states. When solved the

resulting matrix can be used to calculate the effect of an rf pulse on the four level system:

$$R_{\mathbb{CB}} = \begin{pmatrix} e^{\frac{i\tau\omega}{2}} \left(\cos \left[\frac{a\tau}{2} \right] + i\theta_c \sin \left[\frac{a\tau}{2} \right] \right) & ie^{\frac{1}{2}i(2\phi+(2t+\tau)\omega)}\theta_s \sin \left[\frac{a\tau}{2} \right] & 0 & 0 \\ ie^{-\frac{1}{2}i(2\phi+(2t+\tau)\omega)}\theta_s \sin \left[\frac{a\tau}{2} \right] & e^{-\frac{1}{2}i\tau\omega} \left(\cos \left[\frac{a\tau}{2} \right] + i\theta_c \sin \left[\frac{a\tau}{2} \right] \right) & 0 & 0 \\ 0 & 0 & e^{-\frac{1}{2}i\tau\Omega} & 0 \\ 0 & 0 & 0 & e^{-\frac{1}{2}i\tau\Omega} \end{pmatrix}, \quad (\text{B.15})$$

where $\delta = \omega_0 - \omega_{\text{rf}}$, $a = \sqrt{b^2 - \delta^2}$ is the generalised Rabi frequency, $\theta_c = \frac{\delta}{a}$, $\theta_s = \frac{2b}{a}$ and τ is the rf pulse length. The \mathbb{CB} subscript indicates that the rotation is to be applied to a state of the cartesian basis.

We can now consider the YbF interferometer. Several useful results can be drawn from this matrix. Firstly, setting the rf to be resonant with the transition ($\delta = \omega - \omega_0$), a π -pulse can be delivered with the interaction strength b tuned to $b = \pi/2\tau$. We use a π -pulse rather than a $\pi/2$ -pulse because we are interested in creating a superposition sensitive to the EDM, between the $F = 1$ magnetic sublevels. The rf interaction matrix can then be then simplified to

$$R_{\pi\mathbb{CB}} = \begin{pmatrix} 0 & e^{\frac{1}{2}i\tau\Omega} & 0 & 0 \\ -ie^{-\frac{1}{2}i\tau\Omega} & 0 & 0 & 0 \\ 0 & 0 & e^{-\frac{1}{2}i\tau\Omega} & 0 \\ 0 & 0 & 0 & e^{-\frac{1}{2}i\tau\Omega} \end{pmatrix}. \quad (\text{B.16})$$

Application of this to a molecule in the $|0\rangle$ state drives a transition to $|1\pm 1\rangle$, or $|y\rangle$ in the cartesian basis

$$\langle y|\Pi|0\rangle = \langle 1, 1|\mathbb{CB} R_{\pi\mathbb{CB}} \mathbb{CB}|0\rangle + \langle 1, -1|\mathbb{CB} R_{\pi\mathbb{CB}} \mathbb{CB}|0\rangle = 1. \quad (\text{B.17})$$

If the detuning δ is not fixed to resonance, then the rf lineshape can be derived:

$$\langle 1, \pm 1|\mathbb{CB} R_{\pi\mathbb{CB}} \mathbb{CB}|0\rangle = \frac{\pi^2 \sin^2 \left[\frac{1}{2} \sqrt{\pi^2 + \delta^2} \tau^2 \right]}{\pi^2 + \delta^2 \tau^2}. \quad (\text{B.18})$$

Appendix C

YbF Stark shift

The EDM experiment uses the ground state $F = 0$ and $F = 1$, $m_F \pm 1$ states to create a molecular interferometer. Although once in a coherent superposition the interferometer is insensitive to Stark shift detuning in order to perform π pulses we must drive the transition $|F = 0\rangle \rightarrow |F = 1, m_F \pm 1\rangle$. The energy difference between these states does depend on the applied electric field.

In addition electric field mapping experiments rely on theoretical calculations ?? or previous measurements to convert a transition frequency into a stark shift, and in turn to an implied electric field.

Figure C.1 shows the Stark shift energy, expressed as a frequency, between the $|F = 0\rangle \leftrightarrow |F = 1, m_F \pm 1\rangle$ levels. The detailed calculation of this curve may be found in [43]. Assuming the transition linewidth to be large compared to any Zeeman splitting, this is the frequency used to drive rf transitions. The electric field is nominally calculated as $E = V/d$, where V is the applied potential and d is the plate spacing (1.2 cm).




Figure C.1: Stark shift in the $|F = 0\rangle \leftrightarrow |F = 1, m_F \pm 1\rangle$ levels of YbF.

Appendix D

Experiment operation guide EDM Data

This section outlines a checklist, as of May 2007, for operation of the EDM experiment.

D.1 Preparation of the experiment

Starting up

dye laser (380D)	turn on dye laser cooling
	cover dye jet
	dye circulator on
	check pump beam blocked
Ar ⁺ (2580) pump laser	cooling water valves open
	cooling water on
	Ar ⁺ on
	Ar ⁺ set power

Wait 10 minutes before removing dye jet cover. Other equipment may be turned on now. This allows the rf and HV power supplies electronics settle.

Starting up (*cont...*)

rf apparatus	check the rf terminator is in place (if required)
	check transmission line / loops are connected up appropriately
	turn on rf equipment
electric field power supplies	check electric field power supplies properly connected
	check the electric field control is set to off
	turn on supplies

The electric field supplies can be set to the required voltage at this point, allowing the leakage current to be checked. Also, run the conditioner now if required.

Starting up (*cont...*)

ablation laser	turn on Nd-YAG cooling, then turn on the laser
probe laser	remove dye jet cover
	pump the dye laser
	leave one hour to stabilise.

In the dye laser does not lase within 20 minutes minor tuning may be required. Also, the pump beam steering should be checked.

Assuming the dye laser is lasing, leave for 30 minutes, then tune coarsely to maximise output power. Leave a further 20 minutes before proceeding

Starting up (*cont...*)

dye laser	tune up power. 200 mW
probe laser	tune up cavity fringes and iodine spectrometer
	lock the stable lock, then slave to the iodine lock
general	turn on the Ar SF ₆ mix
	turn on the PMT
	check Nd-YAG safety cloth
	block pump beam
software	check probe alignment and power
	load appropriate <i>Scanmaster</i> profile - <i>Align</i>
	start acquisition

At this point a molecular beam should have been observed. The molecular beam now needs to be optimised, as described in appendix A.1.

Starting up (*cont...*)

source	optimise target parameters using appendix A.1
general	degauss outer and inner magnetic shields
	start phaselock
	set electric field voltage
rf	set rf centre-times based upon magnetic field uniformity
	scan frequency over transitions (\sim -5dBm power) to locate centre
	scan power for each transition
	set power and frequency
	save <i>Scanmaster</i> profile

At this stage the rf is optimised; proceed to interference experiments.

Interference scan

connect to the coils to the B-scan box
set scan range is -5 to 5 mA
scan interference curve
note the B-Shift. Fit in *Mathematica* if required
connect coils to B - stepping box
set B-Shift bias
make note of electric field voltage in each state

D.2 Taking EDM Data

Once satisfactory interference scans have been performed, the experiment is ready to take EDM data. The follow procedure must be followed for every cluster.

Acquire EDM cluster

(close previous *Mathematica kernel*)
ensure oscilloscope and degausser are disconnected
note leakage (steady state)
ensure that the machine states are clearly noted (or at least may be inferred from previous block)
initialise using *BOO* script
unplug the oscilloscope
check B shift ok, continue if possible, otherwise correct B.
try to acquire > 10 blocks per clusters

The following checks should be made whilst running:

Continue to monitor

sufficient sensitivity to realise desired measurement

check B -Shift $< 10\%$ CAL

phaselock active

check leakage regularly

Appendix E

Photographs of the EDM experiment




Figure E.1: *a*: View into the top of the experiment, with uppermost section of assembly (including the PMT and LIF detectors) removed. The field plates, mounting assembly and the uppermost rf loop are clearly visible. All bolts inside the shield are aluminium (anodised red), to reduce the possibility of introducing any ferromagnetic material.

b: The finished electric field plates, after electropolishing and gold coating.

c: The electric field plate assembly ready to be installed.

Bibliography

- [1] Edward A. Hinds and P. G. H. Sandars. Experiment to search for P - and T -violating interactions in the hyperfine structure of thallium fluoride. *Phys. Rev. A*, 21(2):480–487, Feb 1980.
- [2] C. A. Baker. Improved experimental limit on the electric dipole moment of the neutron. *Phys. Rev. Lett.*, 97:131801, 2006.
- [3] B. C. Regan, Eugene D. Commins, Christian J. Schmidt, and David DeMille. New limit on the electron electric dipole moment. *Phys. Rev. Lett.*, 88(7):071805, Feb 2002.
- [4] Ronald Steven McNabb. *An Improved Limit on the Electric Dipole Moment of the Muon*. PhD thesis, University of Minnesota, 2003.
- [5] G. W. Bennett, B. Bousquet, H. N. Brown, G. Bunce, R. M. Carey, P. Cushman, G. T. Danby, P. T. Debevec, M. Deile, H. Deng, W. Deninger, S. K. Dhawan, V. P. Druzhinin, L. Duong, E. Efsthadiadis, F. J. M. Farley, G. V. Fedotovitch, S. Giron, F. E. Gray, D. Grigoriev, M. Grosse-Perdekamp, A. Grossmann, M. F. Hare, D. W. Hertzog, X. Huang, V. W. Hughes, and M. Iwasaki. Publisher’s note: Measurement of the positive muon anomalous magnetic moment to 0.7 ppm [phys. rev. lett. 89, 101804 (2002)]. *Phys. Rev. Lett.*, 89(12):129903, Sep 2002.
- [6] Helen R. Quinn. The asymmetry between matter and antimatter. *Physics Today*, February: 30 – 35, 2003.
- [7] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105(4):1413–1415, Feb 1957.
- [8] Richard L. Garwin, Leon M. Lederman, and Marcel Weinrich. Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon. *Phys. Rev.*, 105(4):1415–1417, Feb 1957.
- [9] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π decay of the K_2^0 meson. *Phys. Rev. Lett.*, 13(4):138–140, Jul 1964.

- [10] M.E. Pospelov and I.B. Khriplovich. *Sov. J. Nucl. Phys.*, 53:638, 1991.
- [11] N. Fortson, P. Sandars, and S. Barr. The search for a permanent electric dipole moment. *Phys. Today*, 56N6:33–39, 2003.
- [12] D. F. Nelson, A. A. Schupp, R. W. Pidd, and H. R. Crane. Search for an electric dipole moment of the electron. *Phys. Rev. Lett.*, 2(12):492–495, Jun 1959.
- [13] P.G.H. Sandars. The electric dipole moment of an atom. *Phys. Lett.*, 14:194, 1965.
- [14] Paul Condylis. *An improved measurement of the electron EDM*. PhD thesis, Centre for Cold Matter, Imperial College London, 2004.
- [15] P. G. H. Sandars. Measurability of the proton electric dipole moment. *Phys. Rev. Lett.*, 19(24):1396–1398, Dec 1967.
- [16] Jonathon James Hudson. *Measuring the electric dipole moment of the electron with YbF molecules*. PhD thesis, University of Sussex, 2001.
- [17] J. H. Smith, E. M. Purcell, and N. F. Ramsey. Experimental limit to the electric dipole moment of the neutron. *Phys. Rev.*, 108(1):120–122, Oct 1957.
- [18] Maxim Pospelov and Adam Ritz. Electric dipole moments as probes of new physics. *ANNALS PHYS.*, 318:119, 2005.
- [19] Z. W. Liu and Hugh P. Kelly. Analysis of atomic electric dipole moment in thallium by all-order calculations in many-body perturbation theory. *Phys. Rev. A*, 45(7):R4210–R4213, Apr 1992.
- [20] M. G. Kozlov, A. V. Titov, N. S. Mosyagin, and P. V. Souchko. Enhancement of the electric dipole moment of the electron in the BaF molecule. *Phys. Rev. A*, 56(5):R3326–R3329, Nov 1997.
- [21] H. S. Nataraj, B. K. Sahoo, B. P. Das, and D. Mukherjee. Intrinsic electric dipole moments of paramagnetic atoms: Rubidium and cesium. *Physical Review Letters*, 101:033002, 2008.
- [22] T. M. R. Byrnes, V. A. Dzuba, V. V. Flambaum, and D. W. Murray. Enhancement factor for the electron electric dipole moment in francium and gold atoms. *Physical Review A*, 59:3082, 1999.
- [23] Imperial College Ben Sauer, Centre for cold Matter. Private communication.
- [24] Eugene D. Commins, Stephen B. Ross, David DeMille, and B. C. Regan. Improved experimental limit on the electric dipole moment of the electron. *Phys. Rev. A*, 50(4):2960–2977, Oct 1994.

- [25] Flambaum et al. Fundamental symmetries studies with cold trapped francium atoms at ISAC. *Hyperfine Interactions*, 172:45–51, September 2006.
- [26] M. V. Romalis, W. C. Griffith, J. P. Jacobs, and E. N. Fortson. New limit on the permanent electric dipole moment of ^{199}Hg . *Phys. Rev. Lett.*, 86(12):2505–2508, Mar 2001.
- [27] V. F. Dmitriev and R. A. Sen'kov. Schiff moment of the mercury nucleus and the proton dipole moment. *Physical Review Letters*, 91:212303, 2003.
- [28] W.C. Griffith, M.D. Swallows, E.N. Fortson, and M.V. Romalis. A four vapor cell search for the permanent electric dipole moment of ^{199}Hg . *APS Meeting Abstracts*, pages C4001+, May 2002.
- [29] D. Cho, K. Sangster, and E. A. Hinds. Tenfold improvement of limits on T-violation in thallium fluoride. *Phys. Rev. Lett.*, 63(23):2559–2562, Dec 1989.
- [30] E A Hinds. Testing time reversal symmetry using molecules. *Physica Scripta*, T70:34–41, 1997.
- [31] V. V. Flambaum, I. B. Khriplovich, and O. P. Sushkov. On the possibility to study P- and T- odd nuclear properties in atomic and molecular experiments. *Sov. Phys. JETP*, 60:873, 1984.
- [32] Mikhail Kozlov. Enhancement of the electric dipole moment of the electron in the YbF molecule. *J.PHYS.B*, 30:L607, 1997.
- [33] Y.Y. Dmitriev et al. Calculation of the spin-rotational hamiltonian including p- and p, t-odd weak interaction terms for HgF and PbF molecules. *Phys. Lett. A*, 167:280, 1992.
- [34] Edmund R. Meyer, John L. Bohn, and Michael P. Deskevich. Candidate molecular ions for an electron electric dipole moment experiment. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 73(6):062108, 2006.
- [35] P. G. H. Sandars and E. Lipworth. Electric dipole moment of the cesium atom. a new upper limit to the electric dipole moment of the free electron. *Phys. Rev. Lett.*, 13(24):718–720, Dec 1964.
- [36] S. K. Lamoreaux. Additional motional-magnetic-field considerations for electric-dipole-moment experiments. *Phys. Rev. A*, 53(6):R3705–R3708, Jun 1996.
- [37] D. DeMille, F. Bay, S. Bickman, D. Kowall, D. Krause, S. E. Maxwell, and L. R. Hunter. Investigation of PbO as a system for measuring the electric dipole moment of the electron. *Phys. Rev. A*, 61(5):052507, Apr 2000.
- [38] D. Kowall, F. Bay, S. Bickman, Y. Jiang, and D. DeMille. Precision zeeman-stark spectroscopy of the metastable $a(1)\Sigma^+$ state of PbO. *Physical Review Letters*, 92(13):133007, 2004.

- [39] R.P. Stutz and E.A. Cornell. . *Bull. Am. Soc. Phys*, 89:76, 2004.
- [40] Shafer-Ray. Possibility of 0- g -factor paramagnetic molecules for measurement of the electron’s electric dipole moment. *Phys. Rev. A*, 73(3):034102–+, March 2006.
- [41] J. J. Hudson, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds. Measurement of the electron electric dipole moment using YbF molecules. *Phys. Rev. Lett.*, 89(2):023003, Jun 2002.
- [42] D. DeMille, F. Bay, S. Bickman, D. Kawall, D. Krause, S. E. Maxwell, and L. R. Hunter. Investigation of PbO as a system for measuring the electric dipole moment of the electron. *Phys. Rev. A*, 61(5):052507, Apr 2000.
- [43] B. E. Sauer, Jun Wang, and E. A. Hinds. Laser-rf double resonance spectroscopy of ^{174}YbF in the $X^2 \Sigma^+$ state: Spin-rotation, hyperfine interactions, and the electric dipole moment. *The Journal of Chemical Physics*, 105(17):7412–7420, 1996.
- [44] Giles Redgrave. *Spin spectroscopy of YbF using molecular beam interferometry*. PhD thesis, University of Sussex, 1998.
- [45] M A Player G E Harrison and P G H Sandars. A multichannel phase-sensitive detection method using orthogonal square waveforms. *Journal of Physics E: Scientific Instruments*, 4(10):750–754, 1971.
- [46] B. Efron and R. Tibshirani. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, 1(1):54–75, 1986.
- [47] J. J. Hudson, H. T. Ashworth, D. M. Kara, M. R. Tarbutt, B. E. Sauer, and E. A. Hinds. Pulsed beams as field probes for precision measurement. *ArXiv e-prints*, 705, May 2007.
- [48] S. K. Tokunaga, J. O. Stack, J. J. Hudson, B. E. Sauer, E. A. Hinds, and M. R. Tarbutt. A supersonic beam of cold lithium hydride molecules. *J Chem Phys.*, 126(12):124314, 2006.
- [49] Norman F. Ramsey. *Molecular Beams*. Oxfoed Science Pulications, 1960, 1990.
- [50] M. R. Tarbutt, J. J. Hudson, B. E. Sauer, and E. A. Hinds. Preparation and manipulation of molecules for fundamental physics tests, 2008.
- [51] M. R. Tarbutt, H. L. Bethlem, J. J. Hudson, V. L. Ryabov, V. A. Ryzhov, B. E. Sauer, G. Meijer, and E. A. Hinds. Slowing heavy, ground-state molecules using an alternating gradient decelerator. *Physical Review Letters*, 92(17):173002, 2004.
- [52] D. DeMille et al. High-flux beam source for cold, slow atoms or molecules. *Phys. Rev. Lett.*, 95:173201, 2005.